



US008370448B2

(12) **United States Patent**
Galchev

(10) **Patent No.:** **US 8,370,448 B2**
(45) **Date of Patent:** **Feb. 5, 2013**

(54) **API FOR WORKER NODE RETRIEVAL OF SESSION REQUEST**

(75) Inventor: **Galin Galchev**, Sofia (BG)

(73) Assignee: **SAP AG**, Walldorf (DE)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1087 days.

(21) Appl. No.: **11/025,514**

(22) Filed: **Dec. 28, 2004**

(65) **Prior Publication Data**

US 2006/0143284 A1 Jun. 29, 2006

(51) **Int. Cl.**

G06F 15/167 (2006.01)
G06F 13/00 (2006.01)

(52) **U.S. Cl.** **709/214**; 709/216; 709/213; 719/312; 719/314; 719/318

(58) **Field of Classification Search** 709/213–216, 709/226–229, 244; 719/312, 328–329, 314, 719/318; 718/105

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,951,643	A *	9/1999	Shelton et al.	709/227
6,115,721	A	9/2000	Nagy	
6,167,423	A	12/2000	Chopra et al.	
6,199,179	B1	3/2001	Kauffman et al.	
6,360,303	B1 *	3/2002	Wisler et al.	711/152
6,415,364	B1	7/2002	Bauman et al.	
6,601,089	B1 *	7/2003	Sistare et al.	709/213
6,601,112	B1 *	7/2003	O'Rourke et al.	719/312
6,772,409	B1	8/2004	Chawla et al.	
6,879,995	B1 *	4/2005	Chinta et al.	709/204
7,010,609	B1 *	3/2006	Black et al.	709/230
7,130,891	B2 *	10/2006	Bernardin et al.	709/218
7,188,176	B1 *	3/2007	Nedderman et al.	709/227

7,215,637	B1 *	5/2007	Ferguson et al.	370/230.1
7,246,167	B2 *	7/2007	Kalmuk et al.	709/227
7,254,634	B1 *	8/2007	Davis et al.	709/228
7,373,647	B2 *	5/2008	Kalmuk et al.	719/317
7,395,338	B2	7/2008	Fujinaga	
7,478,390	B2 *	1/2009	Brokenshire et al.	718/100
7,543,067	B2 *	6/2009	Matsubayashi et al.	709/228
7,580,915	B2 *	8/2009	Braun et al.	707/1
7,617,389	B2 *	11/2009	Nikami et al.	712/244
8,140,678	B2 *	3/2012	Fleischer et al.	709/226
2002/0133805	A1	9/2002	Pugh et al.	
2002/0174097	A1	11/2002	Rusch et al.	
2003/0037178	A1 *	2/2003	Vessey et al.	709/319
2003/0058880	A1	3/2003	Sarkinen et al.	

(Continued)

FOREIGN PATENT DOCUMENTS

EP	1027796	6/2004
EP	1027796 B1	6/2004

OTHER PUBLICATIONS

Dagfinn Parnas "SAP Virtual Machine Container", <https://weblogs.sdn.sap.com/pub/wlg/940>, Posted on Oct. 23, 2004, pp. 1-4.

Thomas Smits, "Unbreakable Java—The Java Server that Never Goes Down", *Article*, SAP TechED, Nov. 2004, San Diego and Munich. pp. 1-5.

(Continued)

Primary Examiner — Gregory Todd


(74) *Attorney, Agent, or Firm* — Schwegman Lundberg Woessner P.A.

(57)

ABSTRACT

A method is described that involves, at a worker node, receiving a notification of a request from a queue. The notification contains a handle for a shared memory connection. The queue is implemented with a second shared memory connection. The method involves receiving the request from the shared memory through the connection. The method also involves generating a response to the request at the worker node and sending the response over the shared memory connection.

14 Claims, 13 Drawing Sheets



U.S. PATENT DOCUMENTS

2003/0097360	A1	5/2003	McGuire et al.	
2003/0115190	A1	6/2003	Soderstrom et al.	
2003/0187927	A1*	10/2003	Winchell	709/204
2003/0196136	A1	10/2003	Haynes et al.	
2004/0024881	A1*	2/2004	Elving et al.	709/227
2004/0045014	A1	3/2004	Radhakrishnan	
2004/0167980	A1	8/2004	Doyle et al.	
2004/0215703	A1*	10/2004	Song et al.	709/201
2004/0221294	A1	11/2004	Kalmuk et al.	
2005/0021594	A1	1/2005	Bernardin et al.	
2005/0044197	A1	2/2005	Lai	
2005/0188068	A1	8/2005	Kilian	
2005/0256880	A1	11/2005	Nam Koong et al.	
2005/0262181	A1	11/2005	Schmidt et al.	
2005/0262493	A1	11/2005	Schmidt et al.	
2005/0262512	A1	11/2005	Schmidt et al.	
2006/0053425	A1	3/2006	Berkman et al.	
2006/0064545	A1	3/2006	Wintergerst	
2006/0064549	A1	3/2006	Wintergerst	
2006/0070051	A1	3/2006	Kuck et al.	
2006/0129512	A1	6/2006	Braun et al.	
2006/0129546	A1	6/2006	Braun et al.	
2006/0129981	A1	6/2006	Dostert et al.	
2006/0143284	A1	6/2006	Galchev	
2006/0143328	A1	6/2006	Fleischer	
2006/0143359	A1	6/2006	Dostert	
2006/0143609	A1	6/2006	Stanev	
2006/0143618	A1	6/2006	Fleischer	
2006/0143619	A1	6/2006	Galchev	
2006/0155867	A1	7/2006	Kilian	
2007/0150586	A1	6/2007	Kilian	
2007/0156907	A1	7/2007	Galchev	
2009/0217369	A1	8/2009	Abeni et al.	

OTHER PUBLICATIONS

Norbert Kuck, et al., "SAP VM Container: Using Process Attachable Virtual machines to Provide Isolation and Scalability for Large Servers", *Article*, SAP AG, Walldorf, Germany, 2002 pp. 1-2.

Renee Boucher Ferguson, "SAP Preps New NetWeaver Capabilities", eWeek.com, <http://www.eweek.com/article2/0,1759,1668146,00.asp>, Oct. 7, 2004, pp. 1-6.

SAP NetWeaver Makes Standard Java Robust <http://www.sap.com/company/press/press.epx?pageview=print&pressid=3069>, Oct. 5, 2004, pp. 1-2.

"SAP Beefs Up Java Support Capabilities for New NetWeaver", News Story, Oct. 11, 2004, (Computerworld). <http://www.computerworld.com/printthis/2004/0,4814,96558,00.html>, pp. 1-2.

"SAP Presents New Capabilities for Netweaver", InfoWorld, Oct. 6, 2004 pp. 1-5, http://www.infoworld.com/article/04/10/06/HNsapnetweaver_1.html.

"Load Balancing of the SAP Web as for Java Applications", SAP Library, http://help.sap.com/saphelp_nw04/helpdata/en/b5/e8239ef7bd494f896a84625d4b688f/cont.2 pages, Nov. 2004.

"Failover System", SAP Library, Nov. 2004, http://help.sap.com/saphelp_nw04/helpdata/en/d7/57af0cce55143be5de3a6e4911186.cont.2 pages.

"Http Sessions and Failover of Web Application", SAP Library, Nov. 2004, 1 page. http://help.sap.com/saphelp_nw04/helpdata/en/90/044cc585eaba42b649f16181b0fdf/cont.2

WebLogic RMI Features and Guidelines, eDocs, pp. 1-7, Nov. 2004, http://e-docs.bea.com/wls/docs90/rmi/rmi_api.html.

"Java Technology in SAP Web Application Server", SAP Library, http://help.sap.com/saphelp_nw04/helpdata/en/0d/a3bbeff62847ae10000000a114084/cont Nov. 2004, pp. 1-3.

"High Availability and Failover", SAP Library, http://help.sap.com/sabhelp_nw04/helpdata/en/6c/209da105668c45be12f1d9cc412e2b/cont Nov. 2004, pp. 1-2.

"Failover for Enterprise Beans", SAP Library, http://help.sap.com/saphelp_nw04/helpdata/en/8f/d6e45953a494499ea1b79ab16321d2/cont Nov. 2004, pp. 1-3.

"Failover for Clustered RMI-P4 Remote Objects", SAP Library, http://help.sap.com/saphelp_nw04/helpdata/en/f6/4ef11ab3a52f408aa8a133f33d6cf1/cont Nov. 2004, pp. 1-3.

Grzegorz Czajkowski, "Multitasking without Compromise: a Virtual Machine Evolution", *Paper*, Sun Microsystems Laboratories, 2001, pp. 1-14.

Mick Jordan, et al. "Scaling J2EE Application Servers with the Multi-Tasking Virtual Machine", *Paper*, Sun Microsystems, SMLI TR-2004-135, Jun. 2004, pp. 1-19.

Mick Jordan, et al, "Extending a J2EE Server with Dynamic and Flexible Resource Management", *Paper*, Sun Microsystems, Inc, and School of Computer Science, University of Waterloo, Waterloo, ON, Canada, Oct. 2004, 20 pages.

Grzegorz Czajkowski, et al., "A Multi-User Virtual Machine", *Paper*, Sun Microsystems Laboratories and S³ Lab, Purdue University, West Lafayette, IN, 2003, 14 pages.

Sudhir Movva & Murali Vallath, "Transparent Clustered Database Failover Using JAVA", http://www.quest-pipelines.com/newsletter-v5/0604_A.htm, Apr. 2004, 11 pages.

Diaran Bryce, LogOn Technology Transfer, Kronerg, Germany, Experts' Corner., "Isolates: A New Approach to Multi-Programming in Java Platforms", May 2004, *Article*, 7 pages.

Patrick Tullmann, et al., "Janos: A Java-oriented OS for Active Network Nodes", *Paper*, Flux Research Group, University of Utah, Mar. 2001, 14 pages.

Patrick Doyle, et al., A Modular and Extensible JVM Infrastructure, *Paper*, Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, Ontario, Canada, 14 pages. Jul. 2002.

Sandhya Dwarkadas, et al., "Cashmere-VLM: Remote Memory Paging for Software Distributed Shared Memory", *Paper*, Compaq Cambridge Research Lab, Cambridge, MA and Department of Computer Science, University of Rochester, Rochester, NY, 7 pages, Apr. 1999.

Alex Gontmakher, et al. "Characterizations for Java Memory Behavior", *Paper*, Computer Science Department, Technion, 5 pages, 1997.

"System and Method for Shared Memory/File System Java ObjectGroup Clustered JVM", <http://www.priorartdatabase.com/IPCOM/000021597/>, Published Jan. 2004, 3 pages.

Marcelo Loosco, et al., "A New Distributed Java Virtual Machine for Cluster Computing", Notes in Computer Science, Springer-Verlag, v. 2790, p. 1207-1215, 2003.

D. Dillenbeger, et al., "Building a Java Virtual Machine for Server Applications: The Jvm on Os/390", IBM Systems Journal, vol. 39, No. 1, 2000, pp. 194-210.

Galim Galchev, "Connection Manager for Handling Message Oriented Protocol-Based Requests", U.S. Appl. No. 11/024,498, filed Dec. 28, 2004. Office Action mailed Apr. 1, 2008 (12 pgs).

Rosenberg, David, "Bringing Java to the Enterprise: Oracle on its Java Server Strategy", *IEEE Internet Computing IEEE USA*, vol. 2, No. 2; Database accession No. 5902816, XP002431362; ISSN: 1089-7801, (Mar. 2, 1998), 52-59.

Non-Final Office Action for U.S. Appl. No. 11/322,596, Mailed Jan. 23, 2009, 12 pages.

"U.S. Appl. No. 11/322,596, Non Final Office Action mailed Jan. 23, 2009", 10 pgs.

"U.S. Appl. No. 11/322,596, Response filed Mar. 31, 2009 to Non Final Office Action mailed Jan. 23, 2009", 10 pgs.

Polk, Jennifer, et al., "Oracle Database Net Services Administrator's Guide 10g Release 1 (10.1)", Retrieved on Apr. 26, 2007, reference No. XP002431369, [Online]. Retrieved from the Internet: <URL: http://download-west.oracle.com/docs/cd/B19306_01/network.102/b14212.pdf>, (Oct. 2005), 1-29.

Rosenberg, David, "Bringing Java to the Enterprise: Oracle on its Java Server Strategy", *IEEE Internet Computing IEEE USA*, vol. 2, No. 2, Database accession No. 5902816, XP002431362; ISSN: 1089-7801, (Mar. 2, 1998), 52-59.

"U.S. Appl. No. 11/322,596, Response filed Jan. 20, 2010 to Non Final Office Action mailed Sep. 24, 2009", 8 pgs.

"U.S. Appl. No. 11/322,596 Non-Final Office Action mailed Sep. 24, 2009", 8 pgs.

"U.S. Appl. No. 11/322,596, Final Office Action mailed Jul. 10, 2009", 9 pgs.

"U.S. Appl. No. 11/322,596, Response filed Sep. 10, 2009 to Final Office Action mailed Jul. 10, 2009", 7 pgs.

"U.S. Appl. No. 11/322,596, Final Office Action mailed Apr. 28, 2010", 9 pages.

"U.S. Appl. No. 11/322,596, Response filed May 19, 2010 to Final Office Action mailed Apr. 28, 2010", 8 pgs.

"U.S. Appl. No. 11/322,596, Advisory Action mailed Jun. 1, 2010", 3 pgs.

* cited by examiner




FIG. 1A
(PRIOR ART)




FIG. 1B
(PRIOR ART)




FIG. 2

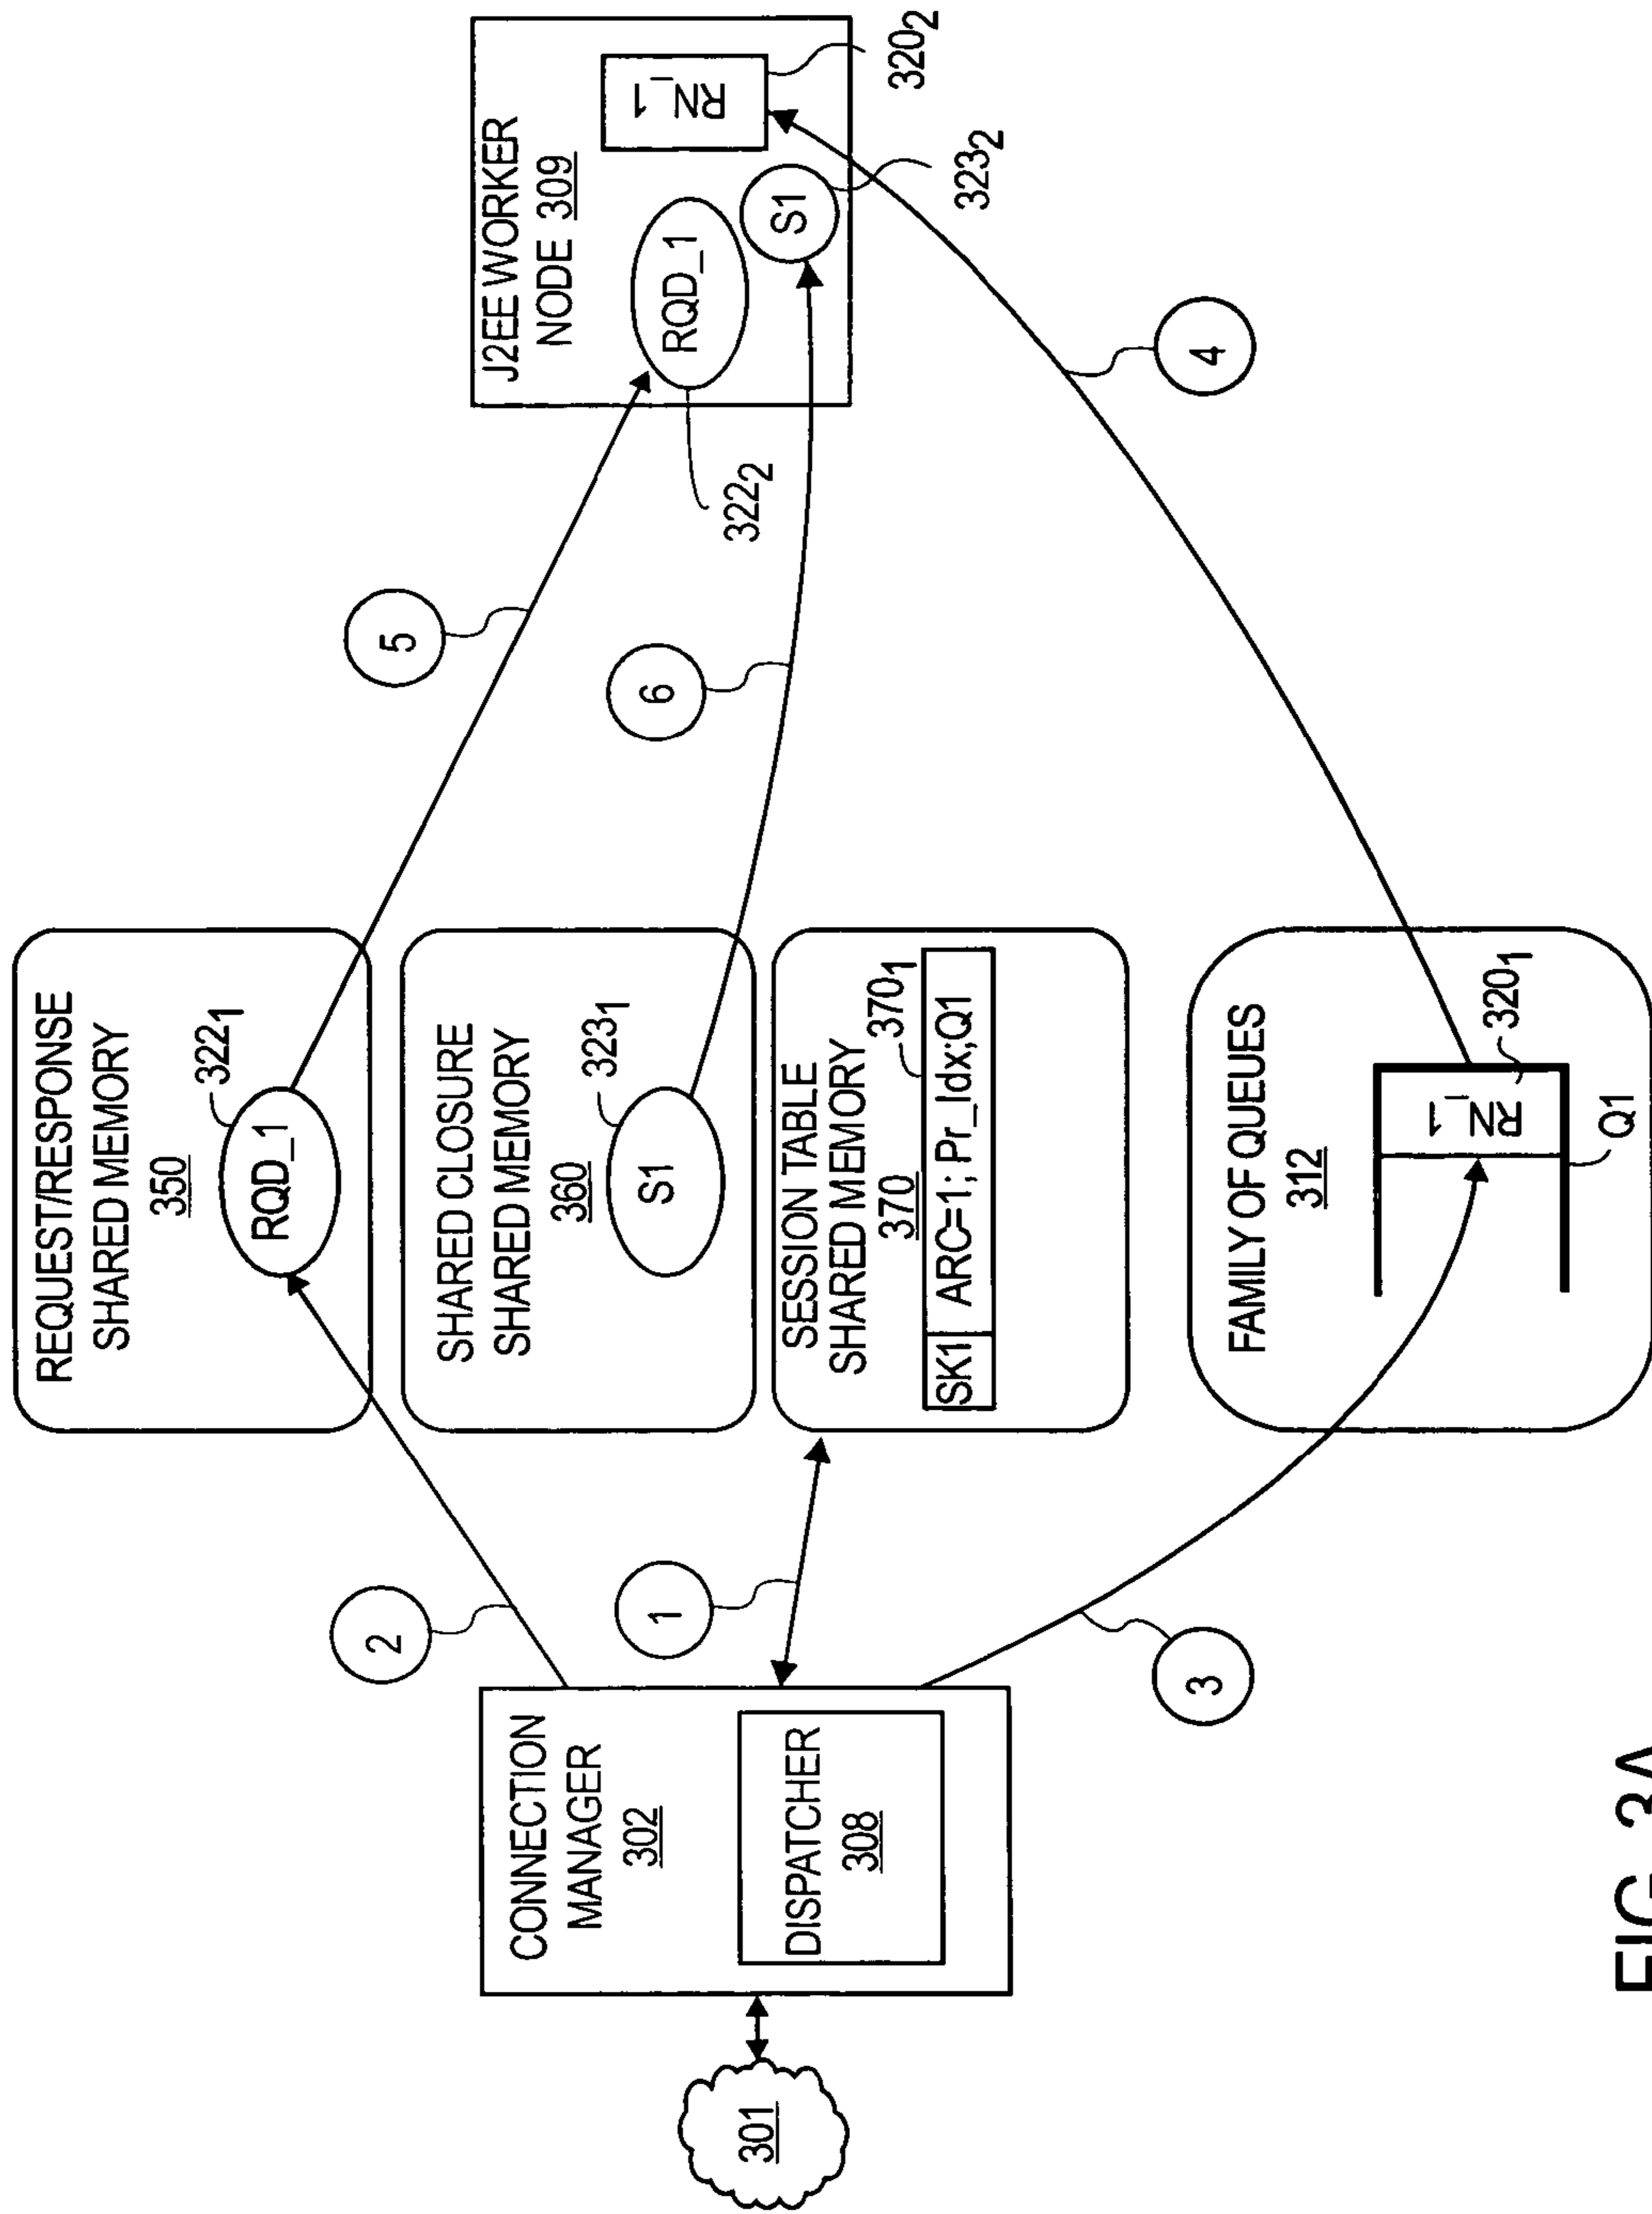


FIG. 3A

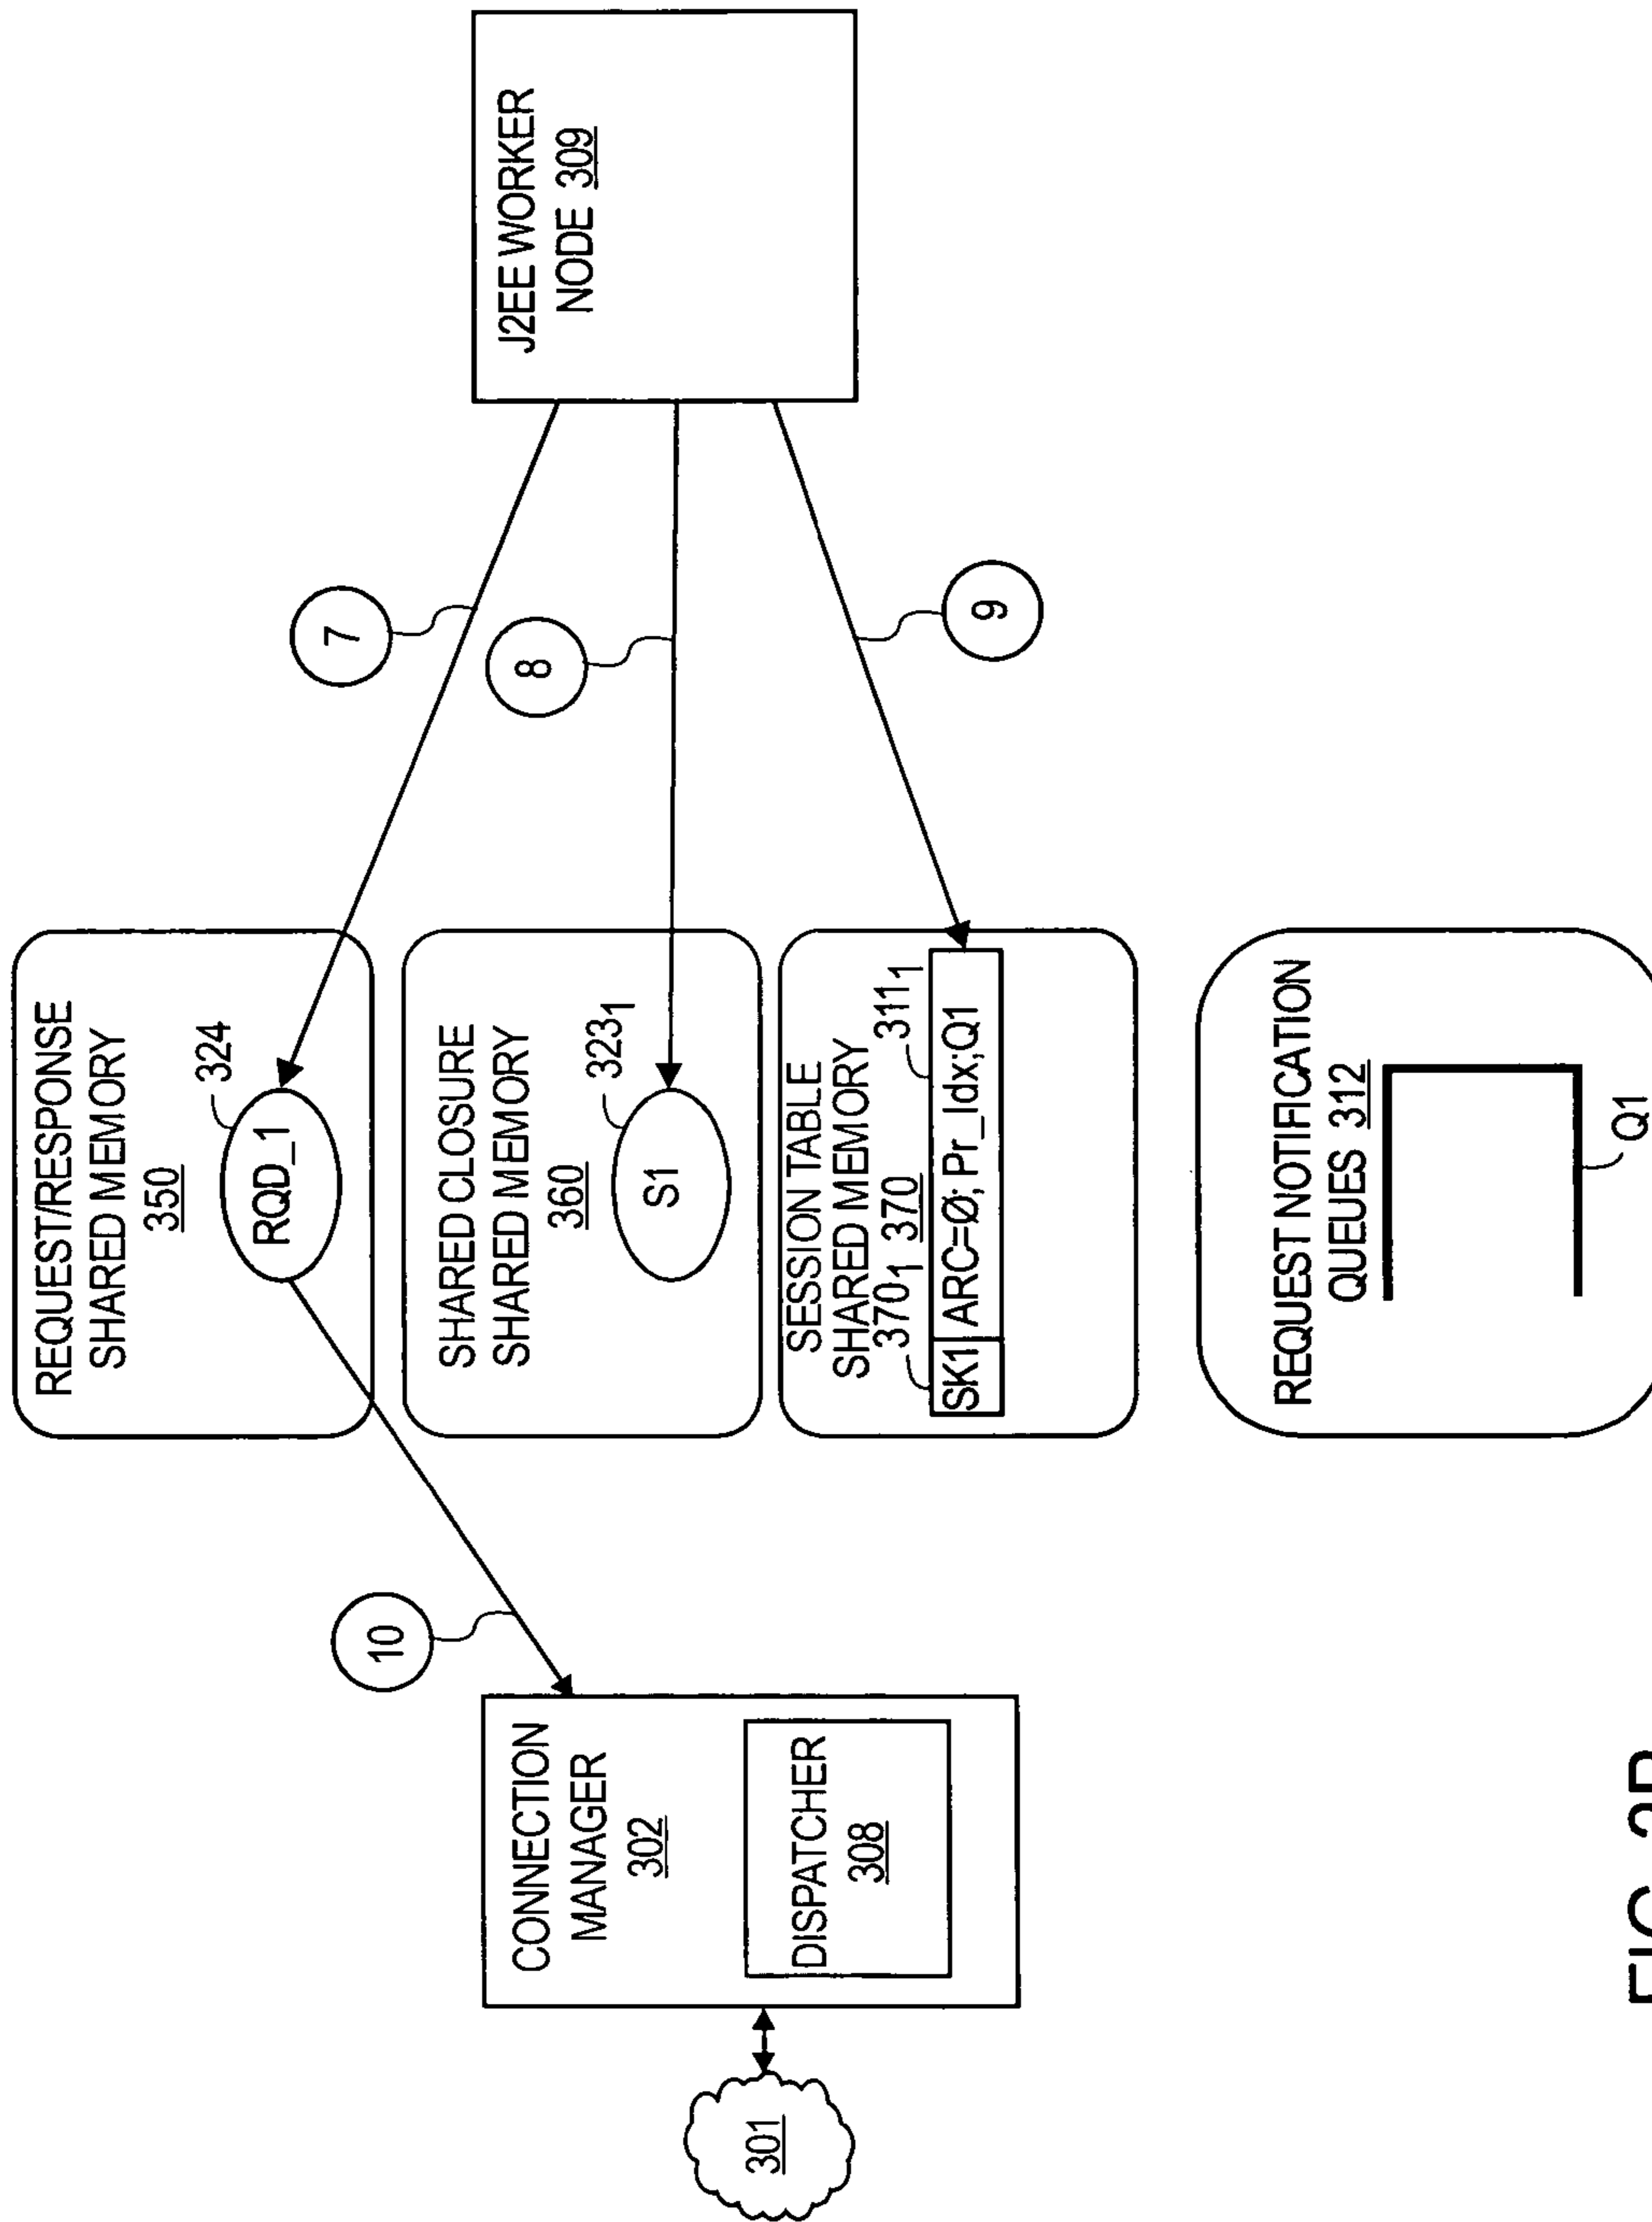


FIG. 3B

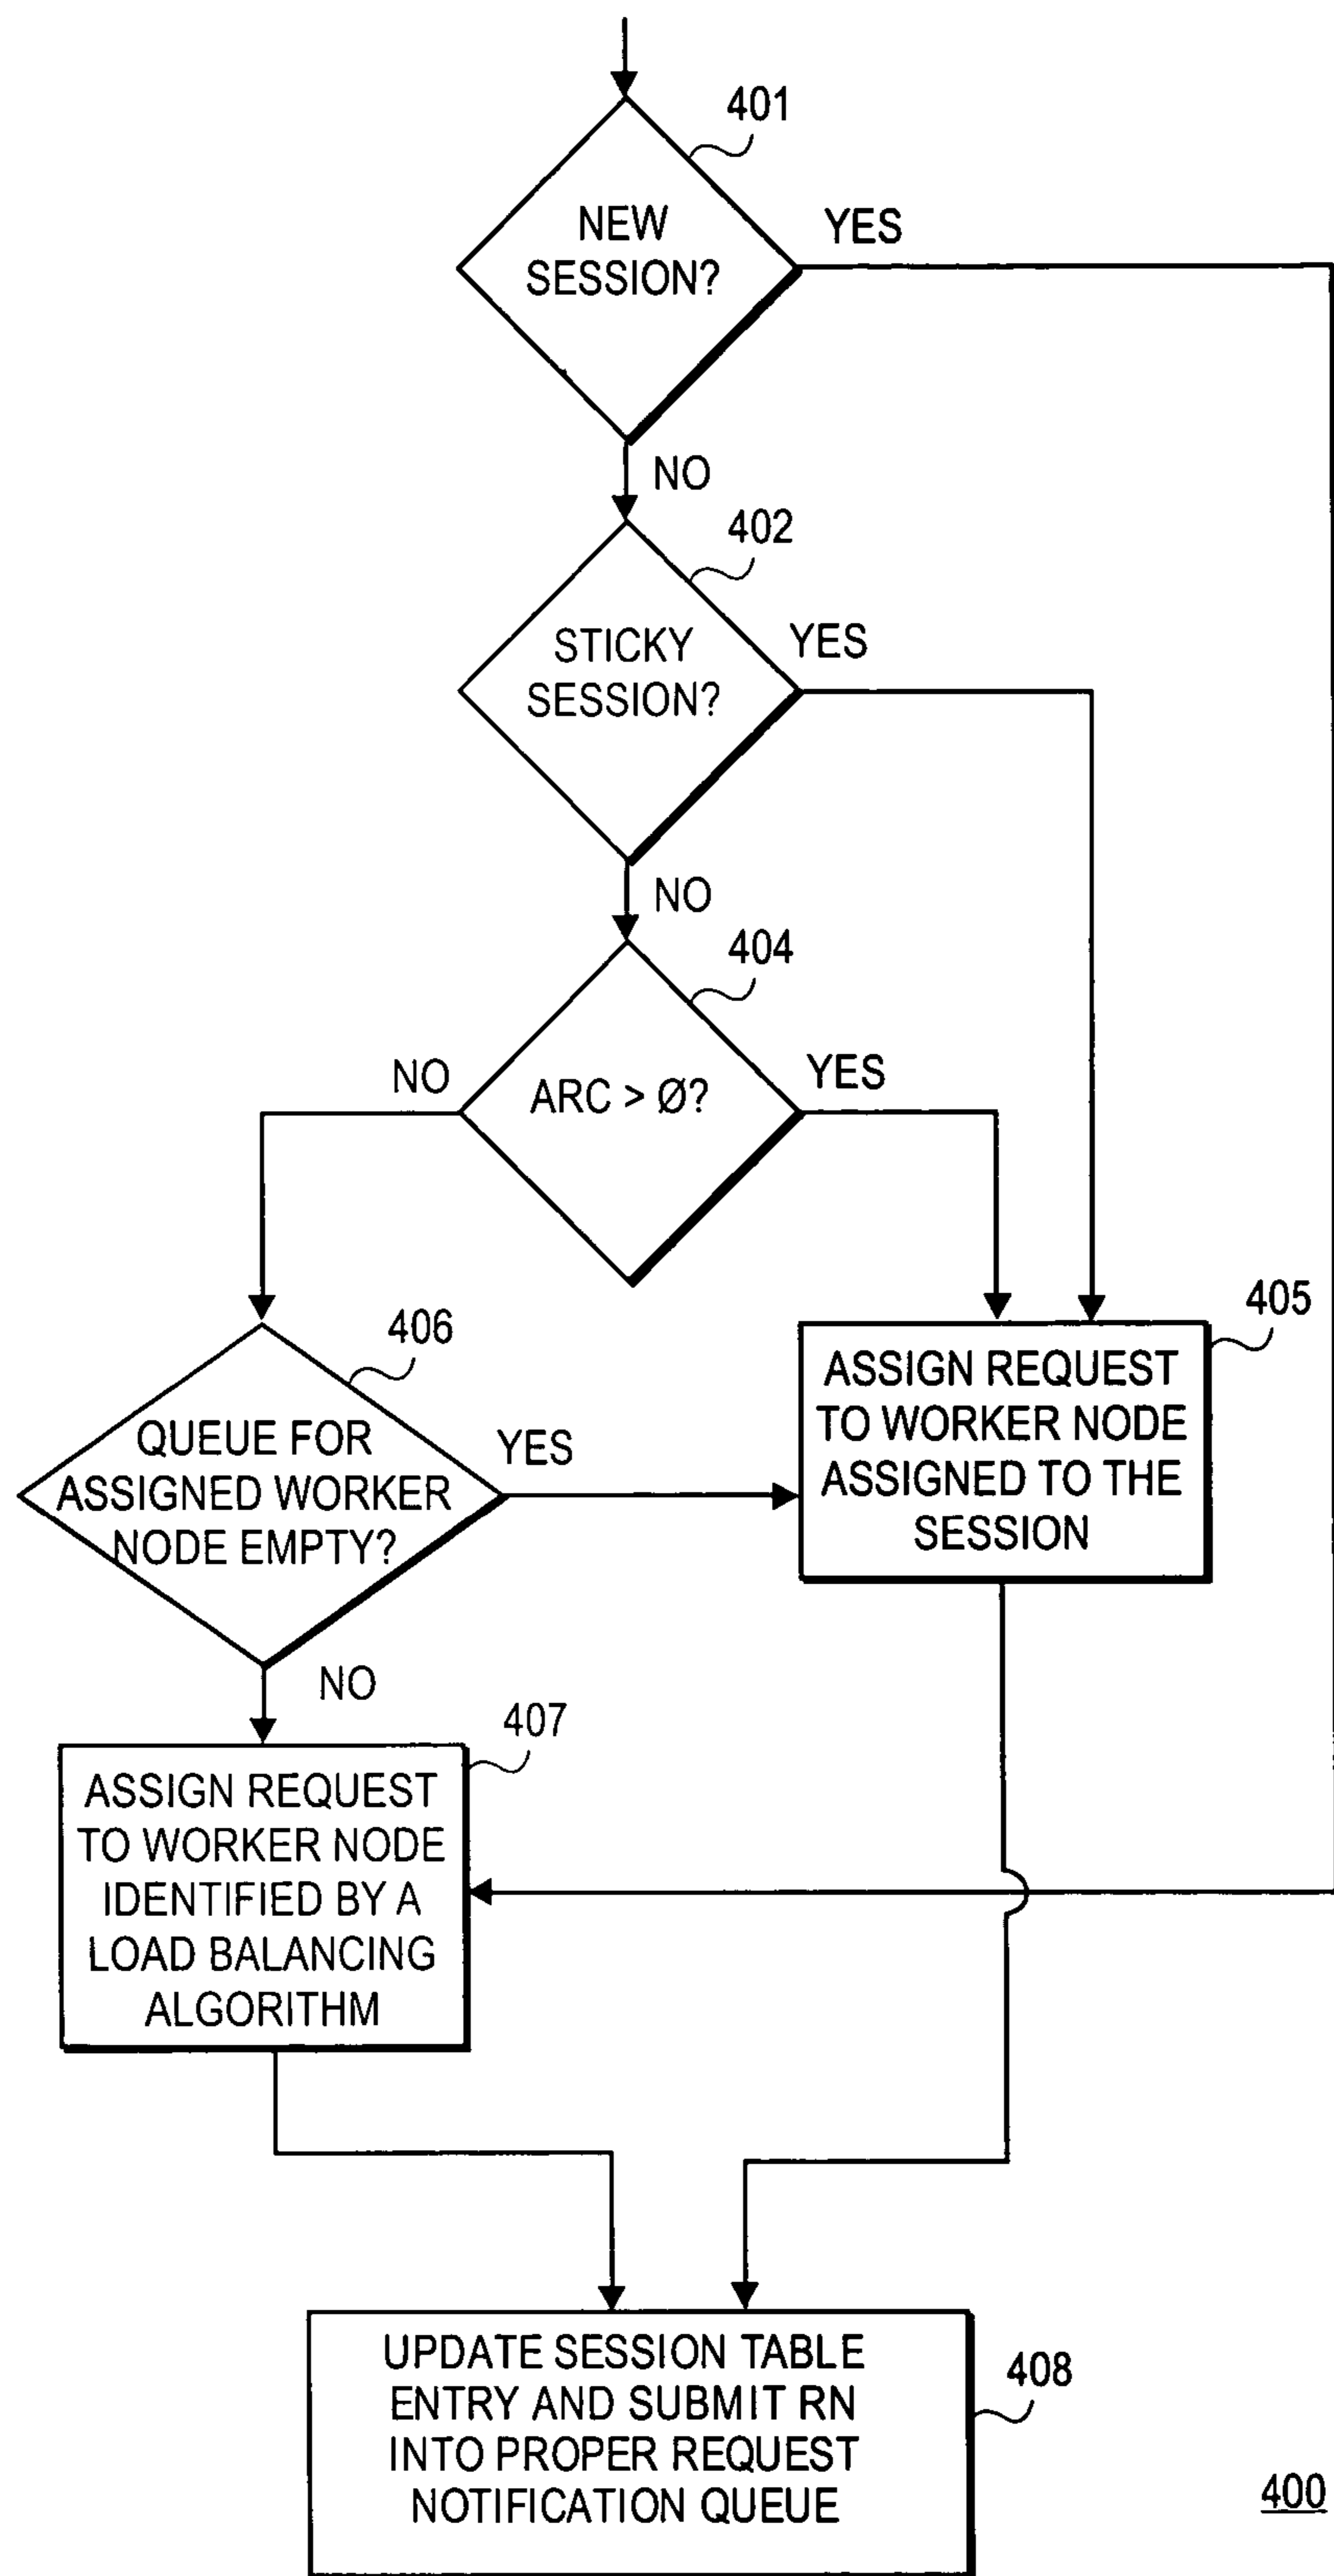


FIG. 4

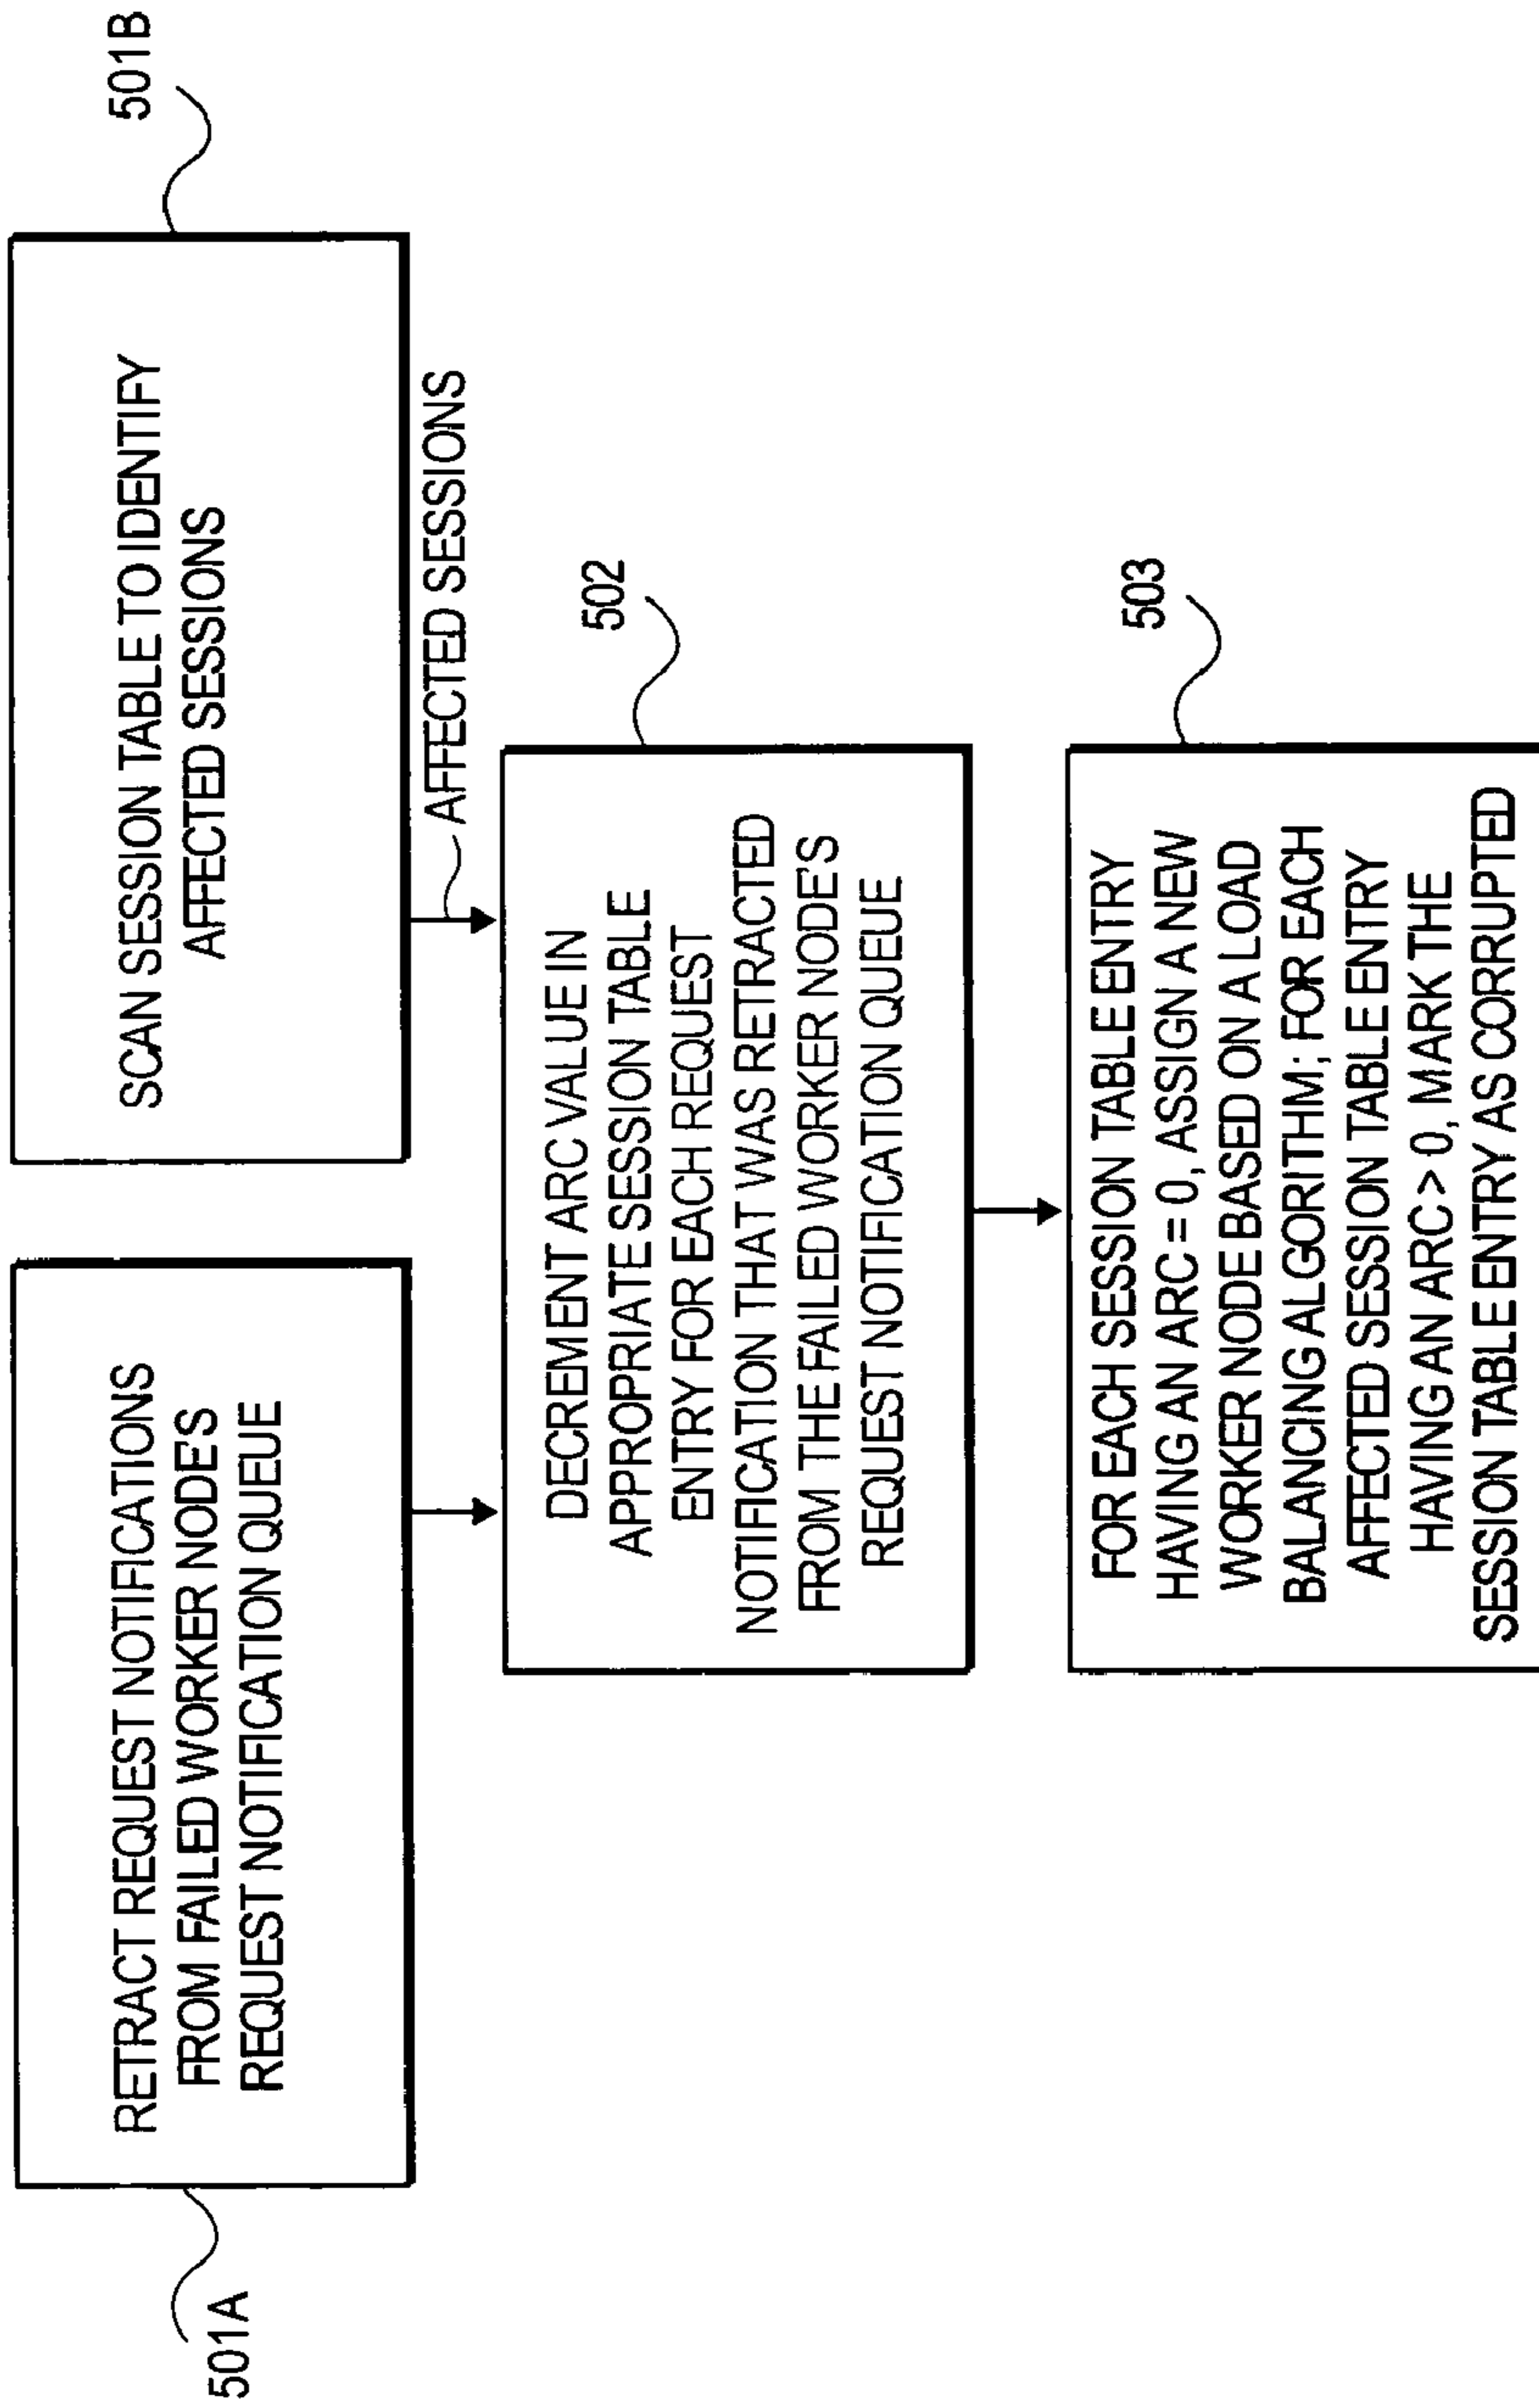


FIG. 5




FIG. 6A




FIG. 6B




FIG. 6C




FIG. 7





FIG. 8



1

API FOR WORKER NODE RETRIEVAL OF SESSION REQUEST

FIELD OF INVENTION

The field of invention relates generally to the software arts; and, more specifically to a architecture that promotes high reliability with multiple worker nodes.

BACKGROUND

Even though standards-based application software (e.g., Java based application software) has the potential to offer true competition at the software supplier level, legacy proprietary software has proven reliability, functionality and integration into customer information systems (IS) infrastructures. Customers are therefore placing operational dependency on standards-based software technologies with caution. Not surprisingly, present day application software servers tend to include both standard and proprietary software suites, and, often, “problems” emerge in the operation of the newer standards-based software, or interoperation and integration of the same with legacy software applications.

The prior art application server **100** depicted in FIGS. **1a, b** provides a good example. FIG. **1a** shows a prior art application server **100** having both an ABAP legacy/proprietary software suite **103** and a Java J2EE standards-based software suite **104**. A connection manager **102** routes requests (e.g., HTTP requests, HTTPS requests) associated with “sessions” between server **100** and numerous clients (not shown in FIG. **1**) conducted over a network **101**. A “session” can be viewed as the back and forth communication over a network **101** between a pair of computing systems (e.g., a particular client and the server).

The back and forth communication typically involves a client (“client”) sending a server **100** (“server”) a “request” that the server **100** interprets into some action to be performed by the server **100**. The server **100** then performs the action and if appropriate returns a “response” to the client (e.g., a result of the action). Often, a session will involve multiple, perhaps many, requests and responses. A single session through its multiple requests may invoke different application software programs.

For each client request that is received by the application server’s connection manager **102**, the connection manager **102** decides to which software suite **103, 104** the request is to be forwarded. If the request is to be forwarded to the proprietary software suite **103**, notification of the request is sent to a proprietary dispatcher **105**, and, the request itself is forwarded into a request/response shared memory **106**. The proprietary dispatcher **105** acts as a load balancer that decides which one of multiple proprietary worker nodes **107₁** through **107_L** are to actually handle the request.

A worker node is a focal point for the performance of work. In the context of an application server that responds to client-server session requests, a worker node is a focal point for executing application software and/or issuing application software code for downloading. The term “working process” generally means an operating system (OS) process that is used for the performance of work and is also understood to be a type of worker node. For convenience, the term “worker node” is used throughout the present discussion.

When a particular proprietary worker node has been identified by dispatcher **105** for handling the aforementioned request, the request is transferred from the request/response shared memory **106** to the identified worker node. The identified worker node processes the request and writes the

2

response to the request into the request/response shared memory **106**. The response is then transferred from the request/response shared memory **106** to the connection manager **102**. The connection manager **102** sends the response to the client via network **101**.

Note that the request/response shared memory **106** is a memory resource that each of worker nodes **107₁** through **107_L** has access to (as such, it is a “shared” memory resource). For any request written into the request/response shared memory **106** by the connection manager **102**, the same request can be retrieved by any of worker nodes **107₁** through **107_L**. Likewise, any of worker nodes **107₁** through **107_L** can write a response into the request/response shared memory **106** that can later be retrieved by the connection manager **102**. Thus the request/response shared memory **106** provides for the efficient transfer of request/response data between the connection manager **102** and the multiple proprietary worker nodes **107₁** through **107_L**.

If the request is to be forwarded to the standards based software suite **104**, notification of the request is sent to the dispatcher **108** that is associated with the standards based software suite **104**. As observed in FIG. **1a**, the standards-based software suite **104** is a Java based software suite (in particular, a Java 2 Enterprise Edition (J2EE) suite) that includes multiple worker nodes **109₁** through **109_N**.

A Java Virtual Machine is associated with each worker node for executing the worker node’s abstract application software code. For each request, dispatcher **108** decides which one of the N worker nodes is best able to handle the request (e.g., through a load balancing algorithm). Because no shared memory structure exists within the standards based software suite **104** for transferring client session information between the connection manager **102** and the worker nodes **109₁** through **109_N**, separate internal connections have to be established to send both notification of the request and the request itself to the dispatcher **108** from connection manager **102** for each worker node. The dispatcher **108** then forwards each request to its proper worker node.

FIG. **1b** shows a more detailed depiction of the J2EE worker nodes **109₁** through **109_N** of the prior art system of FIG. **1a**. Note that each worker node has its own associated virtual machine, and, an extensive amount of concurrent application threads are being executed per virtual machine. Specifically, there are X concurrent application threads (**112₁** through **112_X**) running on virtual machine **113**; there are Y concurrent application threads (**212₁** through **212_Y**) running on virtual machine **213**; . . . and, there are Z concurrent application threads (**N12₁** through **N12_Z**) running on virtual machine **N13**; where, each of X, Y and Z is a large number.

A virtual machine, as is well understood in the art, is an abstract machine that converts (or “interprets”) abstract code into code that is understandable to a particular type of a hardware platform (e.g., a particular type of processor). Because virtual machines operate at the instruction level they tend to have processor-like characteristics, and, therefore, can be viewed as having their own associated memory. The memory used by a functioning virtual machine is typically modeled as being local (or “private”) to the virtual machine. Hence, FIG. **1b** shows local memory **115, 215, . . . N15** allocated for each of virtual machines **113, 213, . . . N13** respectively.

Various problems exist with respect to the prior art application server **100** of FIG. **1a**. To first order, the establishment of connections between the connection manager and the J2EE dispatcher to process a client session adds overhead/inefficiency within the standards based software suite **104**. Moreover, the “crash” of a virtual machine is not an uncommon

event. In the prior art standards suite **104** of FIG. **1a**, requests that are submitted to a worker node for processing are entered into a queue built into the local memory of the virtual machine that is associated with the worker node. If the virtual machine crashes, its in-process as well as its locally queued requests will be lost. As such, potentially, if the requests for a significant number of sessions are queued into the local memory of a virtual machine (e.g., as a direct consequence of the virtual machine's concurrent execution of a significant number of threads), the crash of the virtual machine will cause a significant number of sessions to be "dropped" by the application server **100**.

SUMMARY

A method is described that involves, at a worker node, receiving a notification of a request from a queue. The notification contains a handle for a shared memory connection. The queue is implemented with a second shared memory connection. The method involves receiving the request from the shared memory through the connection. The method also involves generating a response to the request at the worker node and sending the response over the shared memory connection.

FIGURES

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

- FIG. **1a** shows a prior art application server;
- FIG. **1b** shows a more detailed depiction of the J2EE worker nodes of FIG. **1a**;
- FIG. **2** shows an improved application server;
- FIGS. **3a** and **3b** show a session request and response methodology that can be performed by the improved system of FIG. **2**;
- FIG. **4** shows a dispatching methodology;
- FIG. **5** shows a methodology for rescuing sessions that have been targeted for a failed worker node;
- FIGS. **6a** through **6c** depict the rescue of a session whose request notification was targeted for a failed worker node;
- FIG. **7** shows different layers of a shared memory access technology;
- FIG. **8** shows a depiction of a shared closure based shared memory system;
- FIG. **9** shows a depiction of a computing system.

DETAILED DESCRIPTION

1.0 Overview

FIG. **2** shows the architecture of an improved application server that addresses the issues outlined in the Background section.

Comparing FIGS. **1a** and **2**, firstly, note that the role of the connection manager **202** has been enhanced to perform dispatching **208** for the standards based software suite **204** (so as to remove the additional connection overhead associated with the prior art system's standard suite dispatching procedures).

Secondly, the role of a shared memory has been expanded to at least include: a) a first shared memory region **250** that supports request/response data transfers not only for the proprietary suite **203** but also the standards based software suite **204**; b) a second shared memory region **260** that stores session objects having "low level" session state information (i.e., information that pertains to a request's substantive response

such as the identity of a specific servlet invoked through a particular web page); and, c) a third shared memory region **270** that stores "high level" session state information (i.e., information that pertains to the flow management of a request/response pair within the application server (e.g., the number of outstanding active requests for a session)).

Third, request notification queues **212** Q1 through QM, one queue for each of the worker nodes **209₁** through **209_M** has been implemented within the standards-based software suite **204**. As will be described in more detail below, the shared memory structures **250**, **260**, **270** and request notification queues **212** help implement a fast session fail over protection mechanism in which a session that is assigned to a first worker node can be readily transferred to a second worker node upon the failure of the first worker node.

Shared memory is memory whose stored content can be reached by multiple worker nodes. Here, the contents of each of the shared memory regions **250**, **260** and **270** can be reached by each of worker nodes **209₁** through **209_M**. Different types of shared memory technologies may be utilized within the application server **200** and yet still be deemed as being a shared memory structure. For example, shared memory region **250** may be implemented within a "connection" oriented shared memory technology, while shared memory region **260** may be implemented with a "shared closure" oriented shared memory technology. A more thorough discussion of these two different types of shared memory implementations is provided in more detail below in section 5.0 entitled "Implementation Embodiment of Request/Response Shared Memory" and section 6.0 entitled "Implementation Embodiment of Shared Closure Based Shared Memory".

The connection oriented request/response shared memory region **250** effectively implements a transport mechanism for request/response data between the connection manager and the worker nodes. That is, because the connection manager is communicatively coupled to the shared memory, and because shared memory contents can be made accessible to each worker node, the request/response shared memory **250**—at perhaps its broadest level of abstraction—is a mechanism for transporting request/response data between the connection manager and the applicable worker node(s) for normal operation sessions (i.e., no worker node failure) as well as those sessions affected by a worker node crash.

Although the enhancements of the application server **200** of FIG. **2** have been directed to improving the reliability of a combined ABAP/J2EE application server, it is believed that architectural features and methodologies described in more detail further below can be more generally applied to various forms of computing systems that manage communicative sessions, whether or not such computing systems contain different types of application software suites, and whether any such application software suites are standards-based or proprietary. Moreover, it is believed that such architectural features and methodologies are generally applicable irrespective of any particular type of shared memory technology employed.

In operation, the connection manager **202** forwards actual request data to the first shared memory region **250** (request/response shared memory **250**) irregardless of whether the request is to be processed by one of the proprietary worker nodes **207** or one of the standards based worker nodes **204**. Likewise, the connection manager **202** receives response data for a request from the request/response shared memory **250** irregardless if the response was generated by a proprietary worker node or a standards based worker node. With the exception of having to share the request/response shared

memory 250 with the worker nodes 209 of the standards-based software suite 204, the operation of the proprietary suite 203 is essentially the same as that described in the background.

That is, the connection manager 202 forwards request notifications to the proprietary dispatcher 205 and forwards the actual requests to the request/response shared memory 250. The proprietary dispatcher 205 then identifies which one of the proprietary worker nodes 207 is to handle the request. The identified worker node subsequently retrieves the request from the request/response shared memory 250, processes the request and writes the response into the request/response shared memory 250. The response is then forwarded from the request/response shared memory 250 to the connection manager 202 who forwards the response to the client via network 201.

2.0 Processing of a Single Request

FIGS. 3a and 3b show an improved session handling flow that is used within the standards based software suite 204 of the improved application server 200 of FIG. 2. According to this flow, after the connection manager 302 receives a request from network 301 and determines that the request should be handled by the standards-based software suite, the session to which the request belongs is identified (or the request is identified as being the first request of a new session). Here, the connection manager 102 determines the existing session to which the request belongs or that the request is from a new session. through well understood techniques (e.g., through a session identifier found in the header of the received request or a URL path found in the header of the received request).

Then, the dispatcher 308 for the standards-based software suite is invoked. One possible dispatching algorithm that is executed by the dispatcher 308 is described in more detail further below in Section 3.0 entitled “Dispatching Algorithm”. For purposes of the present discussion it is sufficient to realize that the dispatcher 308: 1) accesses and updates at 1 “high level” state information 370₁ for the request’s session in the shared memory session table 370 (hereinafter, referred to as session table 370); 2) determines which one 309 of the M worker nodes should handle the newly arrived request; and 3) submits at 2 the request 322₁ into the request/response shared memory 350 and submits at 3 a request notification 320₁ for the request 322₁ into a request notification queue Q1 that is associated with the worker node 309 identified by the dispatching algorithm. For ease of drawing, FIGS. 3a and 3b only depict the worker node 309 that has been identified by the dispatcher 308.

In an embodiment, there is an entry in the session table 370 for each session being supported by the M worker nodes. If the received request is for a new session (i.e., the received request is the first request of the session), the dispatcher process 308 will create at 1 a new entry 370₁ in the session table 370 for the new session and assign at 2 one of the M worker nodes to handle the session based on a load balancing algorithm. By contrast, if the received request pertains to an already existing session, the dispatcher process 308 will access at 1 the already existing entry 370₁ for the session and use the information therein to effectively determine the proper worker node to handle the request as well as update at 1 the session table entry 370₁. In an embodiment, as will be described in detail further below in Section 3.0, in the case of an already existing session, the determination of the proper worker node may or may not involve the execution of a load balancing algorithm.

In an embodiment, the following items are associated with each session table entry 370₁: 1) a “key” used to access the session table entry 370₁ itself (e.g., session key “SK1”); 2) an active request count (ARC) that identifies the total number of requests for the session that have been received from network 301 but for which a response has not yet been generated by a worker node; 3) an identifier of the worker node 309 that is currently assigned to handle the session’s requests (e.g., “Pr_Idx”, which, in an embodiment, is the index in the process table of the worker node that is currently assigned to handle the session’s requests); and, 4) some form of identification of the request notification queue (Q1) that provides request notifications to the worker node 309 identified in 3) above.

In a further embodiment, each entry in the session table 370 further includes: 1) a flag that identifies the session’s type (e.g., as described in more detail further below in Section 3.0, the flag can indicate a “distributed” session, a “sticky” session, or a “corrupted” session); 2) a timeout value that indicates the maximum amount of time a request can remain outstanding, that is, waiting for a response; 3) the total number of requests that have been received for the session; 4) the time at which the session entry was created; and, 5) the time at which the session entry was last used.

For each request, whether a first request of a new session or a later request for an already established session, the dispatcher’s dispatching algorithm 308 increments the ARC value and at 8 places a “request notification” RN_1 320₁, into the request notification queue Q1 that feeds request notifications to the worker node 309 that is to handle the session. The request notification RN_1 contains both a pointer to the request data RQD_1 322₁ in the request/response shared memory and the session key SK1 in the session table entry for the session.

The pointer is generated by that portion of the connection manager 302 that stores the request data RQD_1 322₁ into the request/response shared memory 350 and is provided to the dispatcher 308. The pointer is used by the worker node 309 to fetch the request data RQD_1 322₁ from the request/response shared memory 350, and, therefore, the term “pointer” should be understood to mean any data structure that can be used to locate and fetch the request data. The session key (or some other data structure in the request notification RN_1 that can be used to access the session table entry 370₁ for the session) is used by the worker node 309 to decrement the ARC counter to indicate the worker node 309 has fully responded to the request.

As will be described in more detail below in section 5.0 entitled “Implementation Embodiment of Request/Response Shared Memory”, according to a particular implementation, the request/response shared memory 350 is connection based. Here, a connection is established between the targeted (assigned) worker node 309 and the connection manager 302 through the request/response shared memory 350 for each request/response cycle that is executed in furtherance of a particular session; and, a handle for a particular connection is used to retrieve a particular request from the request/response shared memory 350 for a particular request/response cycle. According to this implementation, the pointer in the request notification RN is the “handle” for the shared memory 350 connection that is used to fetch request data RQD_1 322₁.

In the case of a first request for a new session, the dispatcher 308 determines which worker node should be assigned to handle the session (e.g., with the assistance of a load balancing algorithm) and places the identity of the worker node’s request notification queue (Q1) into a newly created session table entry 370₁ for the session along with

some form of identification of the worker node itself (e.g., “Pr_Idx”, the index in the process table of the worker node that is currently assigned to handle the session’s requests). For already existing sessions, the dispatcher 308 simply refers to the identify of the request notification queue (Q1) in the session’s session table entry 370₁ in order to understand which request notification queue the request notification RN should be entered into.

In a further embodiment, a single session can entertain multiple “client connections” over its lifespan, where, each client connection corresponds to a discrete time/action period over which the client engages with the server. Different client connections can therefore be setup and torn down between the client and the server over the course of engagement of an entire session. Here, depending on the type of client session, for example in the case of a “distributed” session (described in more detail further below), the dispatcher 308 may decide that a change should be made with respect to the worker node that is assigned to handle the session. If such a change is to be made the dispatcher 308 performs the following within the entry 370₁ for the session: 1) replaces the identity of the “old” worker node with the identity of the “new” worker node (e.g., a “new” Pr_Idx value will replace an “old” Pr_Idx value); and, 2) replaces the identification of the request notification queue for the “old” worker node with an identification of the request notification queue for the “new” worker node.

In another embodiment, over the course a single session and perhaps during the existence of a single client connection, the client may engage with different worker node applications. Here, a different entry in the session table can be entered for each application that is invoked during the session. As such, the level of granularity of a session’s management is drilled further down to each application rather than just the session as a whole. A “session key” (SK1) is therefore generated for each application that is invoked during the session. In an embodiment, the session key has two parts: a first part that identifies the session and a second part that identifies the application (e.g., numerically through a hashing function).

Continuing then with a description of the present example, with the appropriate worker node 309 being identified by the dispatcher 308, the dispatcher 308 concludes with the submission at 2 of the request RQD_1 320₁ into the request/response shared memory 350 and the entry at 3 of a request notification RN_1 320₁ into the queue Q1 that has been established to supply request notifications to worker node 309. The request notification RN_1 320₁ sits in its request notification queue Q1 until the targeted worker node 309 foresees an ability (or has the ability) to process the corresponding request 322₁. Recall that the request notification RN_1 320₁ includes a pointer to the request data itself RQD_1 322₁ as well as a data structure that can be used to access the entry 370₁ in the session table (e.g., the session key SK1).

Comparing FIGS. 2 and 3a, note that with respect to FIG. 2a separate request notification queue is implemented for each worker node (that is, there are M queues, Q1 through QM, for the M worker nodes 209₁ through 209_M, respectively). As will be described in more detail below with respect to FIGS. 5a,b and 6a-c, having a request notification queue for each worker node allows for the “rescue” of a session whose request notification(s) have been entered into the request notification queue of a particular worker node that fails (“crashes”) before the request notification(s) could be serviced from the request notification queue.

When the targeted worker node 309 foresees an ability to process the request 322₁, it looks to its request notification queue Q1 and retrieves at 4 the request notification RN_1

320₁ from the request notification queue Q1. FIG. 3a shows the targeted worker node 309 as having the request notification RN_1 320₂ to reflect the state of the worker node after this retrieval at 4. Recalling that the request notification RN_1 320₁ includes a pointer to the actual request RQD_1 322₁ within the request/response shared memory 350, the targeted worker node 309 subsequently retrieves at 5 the appropriate request RQD_1 322₁ from the request/response shared memory 350. FIG. 3a shows the targeted worker node 309 as having the request RQD_1 322₂ to reflect the state of the worker node after this retrieval at 5. In an embodiment where the request/response shared memory is connection oriented, the pointer to RQD_1 322₁ is a “handle” that the worker node 309 uses to establish a connection with the connection manager 302 and then read at 5 the request RQD_1 322₁ from the request/response shared memory.

The targeted worker node 309 also assumes control of one or more “session” objects S1 323₂ used to persist “low level” session data. Low level session data pertains to the request’s substantive response rather than its routing through the application server. If the request is the first request for a new session, the targeted worker node 309 creates the session object(s) S1 323₂ for the session; or, if the request is a later request of an existing session, the targeted worker node 309 retrieves 6 previously stored session object(s) S1 323₁ from the “shared closure” memory region 360 into the targeted worker node 323₂. The session object(s) S1 may 323₁ be implemented as a number of objects that correspond to a “shared closure”. A discussion of shared closures and an implementation of a shared closure memory region 360 is provided in more detail further below in section 6.0 entitled “Implementation Embodiment of Shared Closure Based Shared Memory”.

With respect to the handling of a new session, the targeted worker node 309 generates a unique identifier for the session object(s) S1 323 according to some scheme. In an embodiment, the scheme involves a random component and an identifier of the targeted worker node itself 309. Moreover, information sufficient to identify a session uniquely (e.g., a sessionid parameter from a cookie that is stored in the client’s browser or the URL path of the request) is found in the header of the request RQD_1 322₂ whether the request is the first request of a new session or a later requests of an existing session. This information can then be used to fetch the proper session object(s) S1 323 for the session.

FIG. 3b depicts the remainder of the session handling process. With the targeted worker node 309 having the request RQD_1 322₂ and low level session state information via session object(s) S1 323₂, the request is processed by the targeted worker node 309 resulting in the production of a response 324 that is to be sent back to the client. The worker node 309 writes at 7 the response 324 into the response/request shared memory 350; and, if a change to the low level session state information was made over the course of generating the response, the worker node 309 writes at 8 updated session object(s) into the shared closure memory 360. Lastly, the worker node 309 decrements at 9 the ARC value in the session table entry 370₁ to reflect the fact that the response process has been fully executed from the worker node’s perspective and that the request has been satisfied. Here, recall that a segment of the request notification RN_1 320₂ (e.g., the session key SK1) can be used to find a “match” to the correct entry 370₁ in the session table 370 in order to decrement of the ARC value for the session.

In reviewing the ARC value across FIGS. 3a and 3b, note that it represents how many requests for the session have been received from network 301 by the connection manager 302

but for which no response has yet been generated by a worker node. In the case of FIG. 3a and 3b only one request is at issue, hence, the ARC value never exceeds a value of 1. Conceivably, multiple requests for the same session could be received from network 301 prior to any responses being generated. In such a case the ARC value will reach a number greater than one that is equal to the number of requests that are queued or are currently being processed by a worker node but for which no response has been generated.

After the response 324 is written at 7 into the request/response shared memory 350, it is retrieved at 10 into the connection manager 302 which then sends it to the client over network 301.

3.0 Dispatching Algorithm

Recall from the discussions of FIGS. 2 and 3a,b that the connection manager 202, 302 includes a dispatcher 208, 308 that executes a dispatching algorithm for requests that are to be processed by any of the M worker nodes 209. FIG. 4 shows an embodiment 400 of a dispatching algorithm that can be executed by the connection manager. The dispatching algorithm 400 of FIG. 4 contemplates the existence of two types of sessions: 1) “distributable”; and, 2) “sticky”.

A distributable session is a session that permits the handling of its requests by different worker nodes over the course of its regular operation (i.e., no worker node crash). A sticky session is a session whose requests are handled by only one worker node over the normal course of its operation. That is, a sticky session “sticks” to the one worker node. According to an implementation, each received request that is to be processed by any of worker nodes 209 is dispatched according to the process 400 of FIG. 4.

Before execution of the dispatching process 400, the connection manager 202, 302 will understand: 1) whether the request is the first request for a new session or is a subsequent request for an already existing session (e.g., in the case of the former, there is no “sessionID” from the client’s browser’s cookie in the header of the request, in the later case there is a such a “sessionID”); and, 2) the type of session associated with the request (e.g., sticky or distributable). In an embodiment, sessions start out as distributable as a default but can be changed to “sticky”, for example, by the worker node that is presently responsible for handling the session.

In the case of a first request for a new session 401, a load balancing algorithm 407 (e.g., round robin based, weight based (e.g., using the number of un-serviced request notifications as weights)) is used to determine which one of the M worker nodes is the proper worker node to handle the request. The dispatching process then writes 408 a new entry for the session into the session table that includes: 1) the sticky or distributable characterization for the session; and, 2) an ARC value of 1 for the session; 3) some form of identification of the worker that has been targeted; and, 4) the request notification queue for the worker node identified by 3). In a further embodiment, a session key is also created for accessing the newly created entry.

If the request is not a first request for a new session 401, whether the received request corresponds to a sticky or distributable session is understood by reference to the session table entry for the session. If the session is a sticky session 402, the request is assigned to the worker node that has been assigned to handle the session 405. According to the embodiment described with respect to FIGS. 3a,b, the identity of the request notification queue (e.g., Q1) for the targeted worker node is listed in the session table entry for the session (note that that the identity of the worker node that is listed in the

session table entry could also be used to identify the correct request notification queue). In a further embodiment, the proper session key is created from information found in the header of the received request.

The ARC value in the session’s session table entry is incremented and the request notification RN for the session is entered into the request notification queue for the worker node assigned to handle the session 408. Recall that the request notification RN includes both a pointer to the request in the request/response shared memory as well as a data structure that can be used by the targeted worker node to access the correct session table entry. The former may be provided by the functionality of the connection manager that stores the request into the request/response shared memory and the later may be the session key.

If the session is a distributable session 402, and if the ARC value obtained from the retrieval of the session’s session table entry is greater than zero 404, the request is assigned to the worker node that has been assigned to handle the session 405. Here, an ARC greater than zero means there still exists at least one previous request for the session for which a response has not yet been generated.

The ARC value for the session is then incremented in the session’s session table entry and the request notification RN for the session is directed to the request notification queue for the worker node assigned to handle the session 408.

If the ARC value is not greater than zero 404, the request is assigned to the worker node that has been assigned to handle the session 405 if the request notification queue for the assigned worker node is empty 406. This action essentially provides an embedded load balancing technique. Since the request notification queue is empty for the worker node that has been assigned to handle the session, the latest request for the session may as well be given to the same worker node.

The ARC value for the session is then incremented in the session’s session table entry and the request notification RN for the session is directed to the request notification queue for the worker node assigned to handle the session 408.

If the ARC value is not greater than zero 404, the request is assigned to a new worker node 407 (through a load balancing algorithm) if the request notification queue for the previously assigned worker node is not empty 406. In this case, there are no un-responded to requests for the session (i.e., ARC=0), the worker node assigned to the session has some backed-up traffic in its request notification queue, and the session is distributable. As such, to improve overall efficiency, the request can be assigned to a new worker node that is less utilized than the previous worker node assigned to handle the session.

The ARC value for the session is incremented in the session’s session table entry and the request notification RN for the session is directed to the request notification queue for the new worker node that has just been assigned to handle the session 408.

4.0 Rescuing Sessions Targeted For a Failed Worker Node

FIGS. 5 and 6a, b, c together describe a scheme for rescuing one or more sessions whose request notifications have been queued into the request notification queue for a particular worker node that crashes before the request notifications are serviced from the request notification queue. FIG. 6a shows an initial condition in which worker nodes 609₁ and 609₂ are both operational. A first request 627 (whose corresponding request notification is request notification 624) for a first session is currently being processed by worker node

609₁. As such, the session object(s) 629 for the first session is also being used by worker node 609₁.

Request notifications 625, 626 are also queued into the request notification queue Q1 for worker node 609₁. Request notification 625 corresponds to a second session that session table 670 entry SK2 and request 628 are associated with. Request notification 626 corresponds to a third session that session table entry SK3 and request 629 are associated with.

FIG. 6b shows activity that transpires after worker node 609₁ crashes at the time of the system state observed in FIG. 6a. Because request notifications 625 and 626 are queued within the queue Q1 for worker node 609₁ at the time of its crash, the second and third sessions are “in jeopardy” because they are currently assigned to a worker node 609₁ that is no longer functioning. Referring to FIGS. 5 and 6b, after worker node 609₁ crashes, each un-serviced request notification 625, 626 is retracted 501a, at 1 from the crashed worker node’s request notification queue Q1; and, each session that is affected by the worker node crash is identified 501b.

Here, recall that in an embodiment, some form of identification of the worker node that is currently assigned to handle a session’s requests is listed in that session’s session table entry. For example, recall that the “Pr_Idx” index value observed in each session table entry in FIG. 6a is an index in the process table of the worker node assigned to handle the request. Assuming the Pr_Idx value has a component that identifies the applicable worker node outright, or can at least be correlated to the applicable worker node, the Pr_Idx values can be used to identify the sessions that are affected by the worker node crash. Specifically, those entries in the session table having a Pr_Idx value that corresponds to the crashed worker are flagged or otherwise identified as being associated with a session that has been “affected” by the worker node crash.

In the particular example of FIG. 6b, the SK1 session table 670 entry will be identified by way of a “match” with the Pr_Idx1 value; the SK2 session table 670 entry will be identified by way of a “match” with the Pr_Idx2 value; and, the SK3 session table 670 entry will be identified by way of a match with the Pr_Idx3 value.

Referring back to FIG. 5 and FIG. 6b, with the retracted request notifications 625, 626 at hand and with the affected sessions being identified, the ARC value is decremented 502, at 2 in the appropriate session table entry for each retracted request notification. Here, recall that each request notification contains an identifier of its corresponding session table entry (e.g., request notification 625 contains session key SK2 and request notification 626 contains session key SK3). Because of this identifier, the proper table entry of decrementing an ARC value can be readily identified.

Thus, the ARC value is decremented for the SK2 session entry in session table 670 and the ARC value is decremented for the SK3 session entry in session table 670. Because the ARC value for each of the SK1, SK2 and SK3 sessions was set equal to 1.0 prior to the crash of worker node 609₁ (referring briefly back to FIG. 6a), the decrement 502, at 2 of the ARC value for the SK2 and SK3 sessions will set the ARC value equal to zero in both of the SK2 and SK3 session table 670 entries as observed in FIG. 6b.

Because the request notification 624 for the SK1 entry had been removed from the request notification queue Q1 prior to the crash, it could not be “retracted” in any way and therefore its corresponding ARC value could not be decremented. As such, the ARC value for the SK1 session remains at 1.0 as observed in FIG. 6b.

Once the decrements have been made for each extracted request notification 502, at 2, decisions can be made as to

which “affected” sessions are salvageable and which “affected” sessions are not salvageable. Specifically, those affected sessions who have decremented down to an ARC value of zero are deemed salvageable; while, those affected sessions who have not decremented down to an ARC value of zero are not deemed salvageable.

Having the ARC value of an affected session decrement down to a value of zero by way of process 502 corresponds to the extraction of a request notification from the failed worker node’s request notification queue for every one of the session’s non-responded to requests. This, in turn, corresponds to confirmation that the requests themselves are still safe in the request/response shared memory 650 and can therefore be subsequently re-routed to another worker node. In the simple example of FIGS. 6a,b, the second SK2 and third SK3 sessions each had an ARC value of 1.0 at the time of the worker node crash, and, each had a pending request notification in queue Q1. As such, the ARC value for the second SK2 and third SK3 sessions each decremented to a value of zero which confirms the existence of requests 628 and 629 in request/response shared memory 650. Therefore the second SK2 and third SK3 sessions can easily be salvaged simply by re-entering request notifications 625 and 626 into the request notification queue for an operational worker node.

The first session SK1 did not decrement down to a value of zero, which, in turn, corresponds to the presence of its request RQD_1 624 being processed by the worker node 609₁ at the time of its crash. As such, the SK1 session will be marked as “corrupted” and eventually dropped.

As another example, assume that each of the request notifications 624, 625, 626 where for the same “first” SK1 session. In this case there would be only one session table 670 entry SK1 in FIG. 6a (i.e., entries SK2 and SK3 would not exist) and the ARC value in entry SK1 would be equal to 3.0 because no responses for any of requests 627, 628 and 629 have yet been generated. The crash of worker node 609₁ and the retraction of all of the request notifications 628, 629 from request notification queue Q1 would result in a final decremented down value of 1.0 for the session. The final ARC value of 1.0 would effectively correspond to the “lost” request 627 that was “in process” by worker node 609₁ at the time of its crash.

Referring to FIGS. 5 and 6b, once the salvageable sessions are known, the retracted request notifications for a same session are assigned to a new worker node based on a load balancing algorithm 503. The retracted request notifications are then submitted to the request notification queue for the new worker node that is assigned to handle the session; and, the corresponding ARC value is incremented in the appropriate session table entry for each re-submitted request notification.

Referring to FIG. 6c, worker node 609₂ is assigned to both the second and third sessions based on the load balancing algorithm. Hence request notifications 625, 626 are drawn being entered at 3 into the request notification queue Q2 for worker node 609₂. The ARC value for both sessions has been incremented back up to a value of 1.0. In the case of multiple retracted request notifications for a same session, in an embodiment, all notifications of the session would be assigned to the same new worker node and submitted to the new worker node’s request notification queue in order to ensure FIFO ordering of the request processing. The ARC value would be incremented once for each request notification.

From the state of the system observed in FIG. 6c, each of request notifications 625, 626 would trigger a set of processes as described in FIGS. 3a,b with worker node 609₂. Impor-

tantly, upon receipt of the request notifications **625**, **626** the new targeted worker node **609**₂ can easily access both the corresponding request data **628**, **629** (through the pointer content of the request notifications and the shared memory architecture) and the session object(s) **622**, **623** (through the request header content and the shared memory architecture).

Note that if different worker nodes were identified as the new target nodes for the second and third sessions, the request notifications **625**, **626** would be entered in different request notification queues.

For distributable sessions, reassignment to a new worker node is a non issue because requests for a distributable session can naturally be assigned to different worker nodes. In order to advocate the implementation of a distributable session, in an implementation, only the session object(s) for a distributable session is kept in shared closure shared memory **660**. Thus, the examples provided above with respect to FIGS. **3a,b** and **6a, b, c** in which low level session object(s) are stored in shared closure shared memory would apply only to distributable sessions. More details concerning shared closure shared memory are provided in section 6.0 "Implementation Embodiment of Shared Closure Shared Memory".

For sticky sessions various approaches exist. According to a first approach, session fail over to a new worker node is not supported and sticky sessions are simply marked as corrupted if the assigned worker node fails (recalling that session table entries may also include a flag that identifies session type).

According to a second approach, session fail over to a new worker node is supported for sticky sessions. According to an extended flavor of this second approach, some sticky sessions may be salvageable while others may not be. According to one such implementation, the session object(s) for a sticky session are kept in the local memory of a virtual machine of the worker node that has been assigned to handle the sticky session (whether the sticky session is rescuable or is not rescuable). Here, upon a crash of a worker node's virtual machine, the session object(s) for the sticky session that are located in the virtual machine's local memory will be lost.

As such, a sticky sessions can be made "rescuable" by configuring it to have its session object(s) serialized and stored to "backend" storage (e.g., to a hard disk file system in the application server or a persisted database) after each request response is generated. Upon a crash of a worker node assigned to handle a "rescuable" sticky session, after the new worker node to handle the sticky session is identified (e.g., through a process such as those explained by FIGS. **5a** and **5b**), the session object(s) for the sticky session are retrieved from backend storage, deserialized and stored into the local memory of the new worker node's virtual machine. Here, sticky sessions that are not configured to have their session object(s) serialized and stored to backend storage after each response is generated are simply lost and will be deemed corrupted.

5.0 Implementation Embodiment of Request/Response Shared Memory

Recall from above that according to a particular implementation, the request/response shared memory **250** has a connection oriented architecture. Here, a connection is established between the targeted worker node and the connection manager across the request/response shared memory **350** for each request/response cycle between the connection manager and a worker node. Moreover, a handle to a particular connection is used to retrieve a particular request from the request/response shared memory.

The connection oriented architecture allows for easy session handling transfer from a crashed worker node to a new worker node because the routing of requests to a new targeted worker node is accomplished merely by routing the handle for a specific request/response shared memory connection to the new worker node. That is, by routing the handle for a request/response shared memory connection to a new worker node, the new worker node can just as easily "connect" with the connection manager to obtain a request as the originally targeted (but now failed) worker node. Here, the "pointer" contained by the request notification is the handle for the request's connection.

FIG. **7** shows an embodiment of an architecture for implementing a connection based queuing architecture. According to the depiction in FIG. **7**, the connection based queuing architecture is implemented at the Fast Channel Architecture (FCA) level **702**. The FCA level **702** is built upon a Memory Pipes technology **701** which is a legacy "semaphore based" request/response shared memory technology **106** referred to in the Background. The FCA level **702** includes an API for establishing connections with the connection manager and transporting requests through them.

In a further embodiment, referring to FIGS. **2** and **7**, the FCA level **702** is also used to implement each of the request notification queues **212**. As such, the request notification queues **212** are also implemented as a shared memory technology. Notably, the handlers for the request notification queues **212** provide more permanent associations with their associated worker nodes. That is, as described, each of the request notification queues **212** is specifically associated with a particular worker node and is "on-going". By contrast, each request/response connection established across request/response shared memory **250** is made easily useable for any worker node (to support fail over to a new worker node), and, according to an implementation, exist only for each request/response cycle.

Above the FCA level **702** is the jFCA level **703**. The jFCA level **703** is essentially an API used by the Java worker nodes and relevant Java parts of the connection manager to access the FCA level **702**. In an embodiment, the jFCA level is modeled after standard Java Networks Socket technology. At the worker node side, however, a "jFCA connection" is created for each separate request/response cycle through request/response shared memory; and, a "jFCA queue" is created for each request notification queue. Thus, whereas a standard Java socket will attach to a specific "port" (e.g., a specific TCP/IP address), according to an implementation, the jFCA API will establish a "jFCA queue" that is configured to implement the request notification queue of the applicable worker node and a "jFCA connection" for each request/response cycle.

Here, an instance of the jFCA API includes the instance of one or more objects to: 1) establish a "jFCA queue" to handle the receipt of request notifications from the worker node's request notification queue; 2) for each request notification, establishing a "jFCA connection" over request/response shared memory with the connection manager so that the corresponding request from the request/response shared memory can be received (through the jFCA's "InputStream"); and, 3) for each received request, the writing of a response back to the same request/response shared memory connection established for the request (through the jFCA's "OutputStream").

In the outbound direction (i.e., from the worker node to the connection manager), in an embodiment, the same jFCA connection that is established through the request/response shared memory between the worker node and the connection

manager for retrieving the request data is used to transport the response back to the connection manager.

In a further embodiment, a service (e.g., an HTTP service) is executed at each worker node that is responsible for managing the flow of requests/responses and the application(s) invoked by the requests sent to the worker node. In a further embodiment, in order to improve session handling capability, the service is provided its own “dedicated thread pool” that is separate from the thread pool that is shared by the worker node’s other applications. By so-doing, a fixed percentage of the worker node’s processing resources are allocated to the service regardless of the service’s actual work load. This permits the service to immediately respond to incoming requests during moments of light actual service work load and guarantees a specific amount of performance under heavy actual service workload.

According to one implementation, each thread in the dedicated thread pool is capable of handling any request for any session. An “available” thread from the dedicated thread pool listens for a request notifications arriving over the jFCA queue. The thread services the request from the jFCA queue and establishes the corresponding jFCA connection with the handler associated with the request notification and reads the request from request/response shared memory. The thread then further handles the request by interacting with the session information associated with the request’s corresponding session.

Each worker node may have its own associated container (s) in which the service runs. A container is used to confine/define the operating environment for the application thread(s) that are executed within the container. In the context of J2EE, containers also provide a family of services that applications executed within the container may use (e.g., Java Naming and Directory Interface (JNDI), Java Database Connectivity (JDBC), Java Messaging Service (JMS) among others).

Different types of containers may exist. For example, a first type of container may contain instances of pages and servlets for executing a web based “presentation” for one or more applications. A second type of container may contain granules of functionality (generically referred to as “components” and, in the context of Java, referred to as “beans”) that reference one another in sequence so that, when executed according to the sequence, a more comprehensive overall “business logic” application is realized (e.g., stringing revenue calculation, expense calculation and tax calculation components together to implement a profit calculation application).

6.0 Implementation Embodiment of Shared Closure Based Shared Memory

Recall from the Background in the discussion pertaining to FIG. 1*b* that the worker nodes 109 depicted therein engage in an extensive number of application threads per virtual machine. FIG. 8 shows worker nodes 809 that can be viewed as a detailed depiction of an implementation for worker nodes 209 of FIG. 2; where, the worker nodes 209, 809 are configured with less application threads per virtual machine than the prior art approach of FIG. 1*b*. Less application threads per virtual machine results in less application thread crashes per virtual machine crash; which, in turn, should result in the new standards-based suite 204 of FIG. 2 exhibiting better reliability than the prior art standards-based suite 104 of FIG. 1*a*.

According to the depiction of FIG. 8, which is an extreme representation of the improved approach, only one application thread exists per virtual machine (specifically, thread 122 is being executed by virtual machine 123; thread 222 is being executed by virtual machine 223; . . . and, thread M22 is being

executed by virtual machine M23). In practice, the worker nodes 809 of FIG. 8 may permit a limited number of threads to be concurrently processed by a single virtual machine rather than only one.

In order to concurrently execute a comparable number of application threads as the prior art worker nodes 109 of FIG. 1*b*, the improved worker nodes 809 of FIG. 8 instantiate more virtual machines than the prior art worker nodes 109 of FIG. 1*b*. That is, $M > N$.

Thus, for example, if the prior art worker nodes 109 of FIG. 1*b* have 10 application threads per virtual machine and 4 virtual machines (e.g., one virtual machine per CPU in a computing system having four CPUs) for a total of $4 \times 10 = 40$ concurrently executed application threads for the worker nodes 109 as a whole, the improved worker nodes 809 of FIG. 8 may only permit a maximum of 5 concurrent application threads per virtual machine and 6 virtual machines (e.g., 1.5 virtual machines per CPU in a four CPU system) to implement a comparable number ($5 \times 6 = 30$) of concurrently executed threads as the prior art worker nodes 109 of FIG. 1*b*.

Here, the prior art worker nodes 109 instantiate one virtual machine per CPU while the improved worker nodes 809 of FIG. 8 can instantiate multiple virtual machines per CPU. For example, in order to achieve 1.5 virtual machines per CPU, a first CPU may be configured to run a single virtual machine while a second CPU in the same system may be configured to run a pair of virtual machines. By repeating this pattern for every pair of CPUs, such CPU pairs will instantiate 3 virtual machines per CPU pair (which corresponds to 1.5 virtual machines per CPU).

Recall from the discussion of FIG. 1*b* that a virtual machine can be associated with its own local memory. Because the improved worker nodes 809 of FIG. 8 instantiate more virtual machines than the prior art working nodes 109 of FIG. 1*b*, in order to conserve memory resources, the virtual machines 123, 223, . . . M23 of the worker nodes 809 of FIG. 8 are configured with less local memory space 125, 225, . . . M25 than the local memory space 115, 215, . . . N15 of virtual machines 113, 213, . . . N23 of FIG. 1*b*. Moreover, the virtual machines 123, 223, . . . M23 of the worker nodes 809 of FIG. 8 are configured to use a shared memory 860. Shared memory 860 is memory space that contains items that can be accessed by more than one virtual machine (and, typically, any virtual machine configured to execute “like” application threads that is coupled to the shared memory 860).

Thus, whereas the prior art worker nodes 109 of FIG. 1*b* use fewer virtual machines with larger local memory resources containing objects that are “private” to the virtual machine; the worker nodes 809 of FIG. 8, by contrast, use more virtual machines with less local memory resources. The less local memory resources allocated per virtual machine is compensated for by allowing each virtual machine to access additional memory resources. However, owing to limits in the amount of available memory space, this additional memory space 860 is made “shareable” amongst the virtual machines 123, 223, . . . M23.

According to an object oriented approach where each of virtual machines 123, 223, . . . M23 does not have visibility into the local memories of the other virtual machines, specific rules are applied that mandate whether or not information is permitted to be stored in shared memory 860. Specifically, to first order, according to an embodiment, an object residing in shared memory 860 should not contain a reference to an object located in a virtual machine’s local memory because an object with a reference to an unreachable object is generally deemed “non useable”.

That is, if an object in shared memory **860** were to have a reference into the local memory of a particular virtual machine, the object is essentially non useable to all other virtual machines; and, if shared memory **860** were to contain an object that was useable to only a single virtual machine, the purpose of the shared memory **860** would essentially be defeated.

In order to uphold the above rule, and in light of the fact that objects frequently contain references to other objects (e.g., to effect a large process by stringing together the processes of individual objects; and/or, to effect relational data structures), “shareable closures” are employed. A “closure” is a group of one or more objects where every reference stemming from an object in the group that references another object does not reference an object outside the group. That is, all the object-to-object references of the group can be viewed as closing upon and/or staying within the confines of the group itself. Note that a single object without any references stemming from can be viewed as meeting the definition of a closure.

If a closure with a non shareable object were to be stored in shared memory **860**, the closure itself would not be shareable with other virtual machines, which, again, defeats the purpose of the shared memory **860**. Thus, in an implementation, in order to keep only shareable objects in shared memory **860** and to prevent a reference from an object in shared memory **860** to an object in a local memory, only “shareable” (or “shared”) closures are stored in shared memory **860**. A “shared closure” is a closure in which each of the closure’s objects are “shareable”.

A shareable object is an object that can be used by other virtual machines that store and retrieve objects from the shared memory **860**. As discussed above, in an embodiment, one aspect of a shareable object is that it does not possess a reference to another object that is located in a virtual machine’s local memory. Other conditions that an object must meet in order to be deemed shareable may also be effected. For example, according to a particular Java embodiment, a shareable object must also possess the following characteristics: 1) it is an instance of a class that is serializable; 2) it is an instance of a class that does not execute any custom serializing or deserializing code; 3) it is an instance of a class whose base classes are all serializable; 4) it is an instance of a class whose member fields are all serializable; 5) it is an instance of a class that does not interfere with proper operation of a garbage collection algorithm; 6) it has no transient fields; and, 7) its finalize () method is not overwritten.

Exceptions to the above criteria are possible if a copy operation used to copy a closure into shared memory **860** (or from shared memory **860** into a local memory) can be shown to be semantically equivalent to serialization and deserialization of the objects in the closure. Examples include instances of the Java 2 Platform, Standard Edition 1.3 java.lang.String class and java.util.Hashtable class.

A container is used to confine/define the operating environment for the application thread(s) that are executed within the container. In the context of J2EE, containers also provide a family of services that applications executed within the container may use (e.g., (e.g., Java Naming and Directory Interface (JNDI), Java Database Connectivity (JDBC), Java Messaging Service (JMS) among others).

Different types of containers may exist. For example, a first type of container may contain instances of pages and servlets for executing a web based “presentation” for one or more applications. A second type of container may contain granules of functionality (generically referred to as “components” and, in the context of Java, referred to as “beans”) that reference one another in sequence so that, when executed accord-

ing to the sequence, a more comprehensive overall “business logic” application is realized (e.g., stringing revenue calculation, expense calculation and tax calculation components together to implement a profit calculation application).

7.0 Additional Comments

The architectures and methodologies discussed above may be implemented with various types of computing systems such as an application server that includes a Java 2 Enterprise Edition (“J2EE”) server that supports Enterprise Java Bean (“EJB”) components and EJB containers (at the business layer) and/or Servlets and Java Server Pages (“JSP”) (at the presentation layer). Of course, other embodiments may be implemented in the context of various different software platforms including, by way of example, Microsoft.NET, Windows/NT, Microsoft Transaction Server (MTS), the Advanced Business Application Programming (“ABAP”) platforms developed by SAP AG and comparable platforms.

Processes taught by the discussion above may be performed with program code such as machine-executable instructions which cause a machine (such as a “virtual machine”, a general-purpose processor disposed on a semiconductor chip or special-purpose processor disposed on a semiconductor chip) to perform certain functions. Alternatively, these functions may be performed by specific hardware components that contain hardwired logic for performing the functions, or by any combination of programmed computer components and custom hardware components.

An article of manufacture may be used to store program code. An article of manufacture that stores program code may be embodied as, but is not limited to, one or more memories (e.g., one or more flash memories, random access memories (static, dynamic or other)), optical disks, CD-ROMs, DVD ROMs, EPROMs, EEPROMs, magnetic or optical cards or other type of machine-readable media suitable for storing electronic instructions. Program code may also be downloaded from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of data signals embodied in a propagation medium (e.g., via a communication link (e.g., a network connection)).

FIG. 9 is a block diagram of a computing system **900** that can execute program code stored by an article of manufacture. It is important to recognize that the computing system block diagram of FIG. 9 is just one of various computing system architectures. The applicable article of manufacture may include one or more fixed components (such as a hard disk drive **902** or memory **905**) and/or various movable components such as a CD ROM **903**, a compact disc, a magnetic tape, etc. In order to execute the program code, typically instructions of the program code are loaded into the Random Access Memory (RAM) **905**; and, the processing core **906** then executes the instructions. The processing core may include one or more processors and a memory controller function. A virtual machine or “interpreter” (e.g., a Java Virtual Machine) may run on top of the processing core (architecturally speaking) in order to convert abstract code (e.g., Java bytecode) into instructions that are understandable to the specific processor(s) of the processing core **906**.

It is believed that processes taught by the discussion above can be practiced within various software environments such as, for example, object-oriented and non-object-oriented programming environments, Java based environments (such as a Java 2 Enterprise Edition (J2EE) environment or environments defined by other releases of the Java standard), or other environments (e.g., a .NET environment, a Windows/NT environment each provided by Microsoft Corporation).

19

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

The invention claimed is:

1. A method, comprising:
 - entering a notification of a request into a request notification queue, said notification including a handle of a connection to a request/response shared memory, said notification not including request data corresponding to said request, said request notification queue being dedicated to storing request notifications for a worker node, said request corresponding to a session;
 - entering said request data corresponding to said request into said request/response shared memory; and
 - based on said entering of said notification of said request and said entering of said request data, incrementing an active request counter in an entry of a session table, said entry corresponding to said session, said session table being stored in a session table shared memory.
2. The method of claim 1 wherein said request/response shared memory is configured to be shared by worker nodes of a standards-based software suite and worker nodes of a proprietary software suite.
3. The method of claim 1 wherein at least one of said entering of said notification of said request and said entering of said request data corresponding to said request includes using an instance of an object of an application program interface.
4. The method of claim 1, further comprising:
 - detecting a failure of said worker node;
 - extracting said notification of said request from said request notification queue;
 - decrementing said active request counter in said entry of said session table; and
 - based on said active request counter in said entry of said session table being equal to zero, submitting said notification of said request into an additional request notification queue, said additional request notification queue being dedicated to storing request notifications for an additional worker node.
5. The method of claim 4, further comprising selecting said additional worker node based on said additional request notification queue being empty and said additional worker node being assigned to handle said session.
6. The method of claim 4, further comprising, based on said active request counter in said entry of said session table being greater than zero, marking said session as corrupted.
7. The method of claim 1, wherein said connection to said request/response shared memory is implemented as an input/output stream across said request/response shared memory.
8. A non-transitory computer-readable storage medium having instructions embodied thereon, the instructions executable by a processor for performing a method, the method comprising:
 - entering a notification of a request into a request notification queue, said notification including a handle of a connection to a request/response shared memory, said notification not include request data corresponding to said request, said request notification queue being dedi-

20

- cated to store request notifications for a worker node, said request corresponding to said session;
 - entering said request data corresponding to said request into said request/response shared memory; and
 - based on said entering of said notification of said request and said entering of said request data, incrementing an active request counter in an entry of a session table, said entry corresponding to said session, said session table being stored in a session table shared memory.
9. The non-transitory computer-readable storage medium of claim 8 wherein said request/response shared memory is configured to be shared by worker nodes of a standards-based software suite and worker nodes of a proprietary software suite.
10. The non-transitory computer-readable storage medium of claim 8 wherein at least one of said entering of said notification of said request and said entering of said request data corresponding to said request includes using an instance of an object of an application program interface.
11. A system comprising:
 - at least one processor; and
 - a worker node implemented by the at least one processor and configured to:
 - retrieve a notification of a request from a request notification queue, the notification of the request including a handle for accessing a request/response shared memory and a data structure for accessing an entry of a shared session table in a session table shared memory, the entry corresponding to a session associated with the request, the notification of the request not including the request data, the request notification queue being dedicated to storing request notifications for the worker node;
 - use the handle to read request data corresponding to the request from the request/response shared memory;
 - process the request based on the request data;
 - use the handle to write response data to the request/response shared memory; and
 - use the data structure to decrement an active request count in the entry of the shared session table in the session table shared memory based on the writing of the response data to the request/response shared memory.
12. The system of claim 11, further comprising foreseeing that the worker node has an ability to process the request and wherein the retrieving of the notification of the request is based on the foreseeing.
13. The system of claim 11, wherein the worker node is further configured to:
 - determine the request corresponds to a distributable session; and
 - based on the determining, keep a session object corresponding to the worker node in a shared closure shared memory, the request/response shared memory being a first region of a shared memory, the shared closure shared memory being a second region of the shared memory, the shared memory being a memory resource accessible by a plurality of worker nodes, the plurality of worker nodes including the worker node.
14. The system of claim 11, wherein the notification of the request further includes a session key and the worker node is further configured to find the entry of the shared session table based on the session key.